AN APPROACH TO SELF-EXTENSION:

REASONING ABOUT DIFFERENT EPISTEMIC ACTIONS AND UNCERTAINTY FOR AUTONOMOUS KNOWLEDGE GATHERING

Marc Hanheide, University of Lincoln

LINCOLN ROBOTICS

DISCLAIMER: NO ICUB!

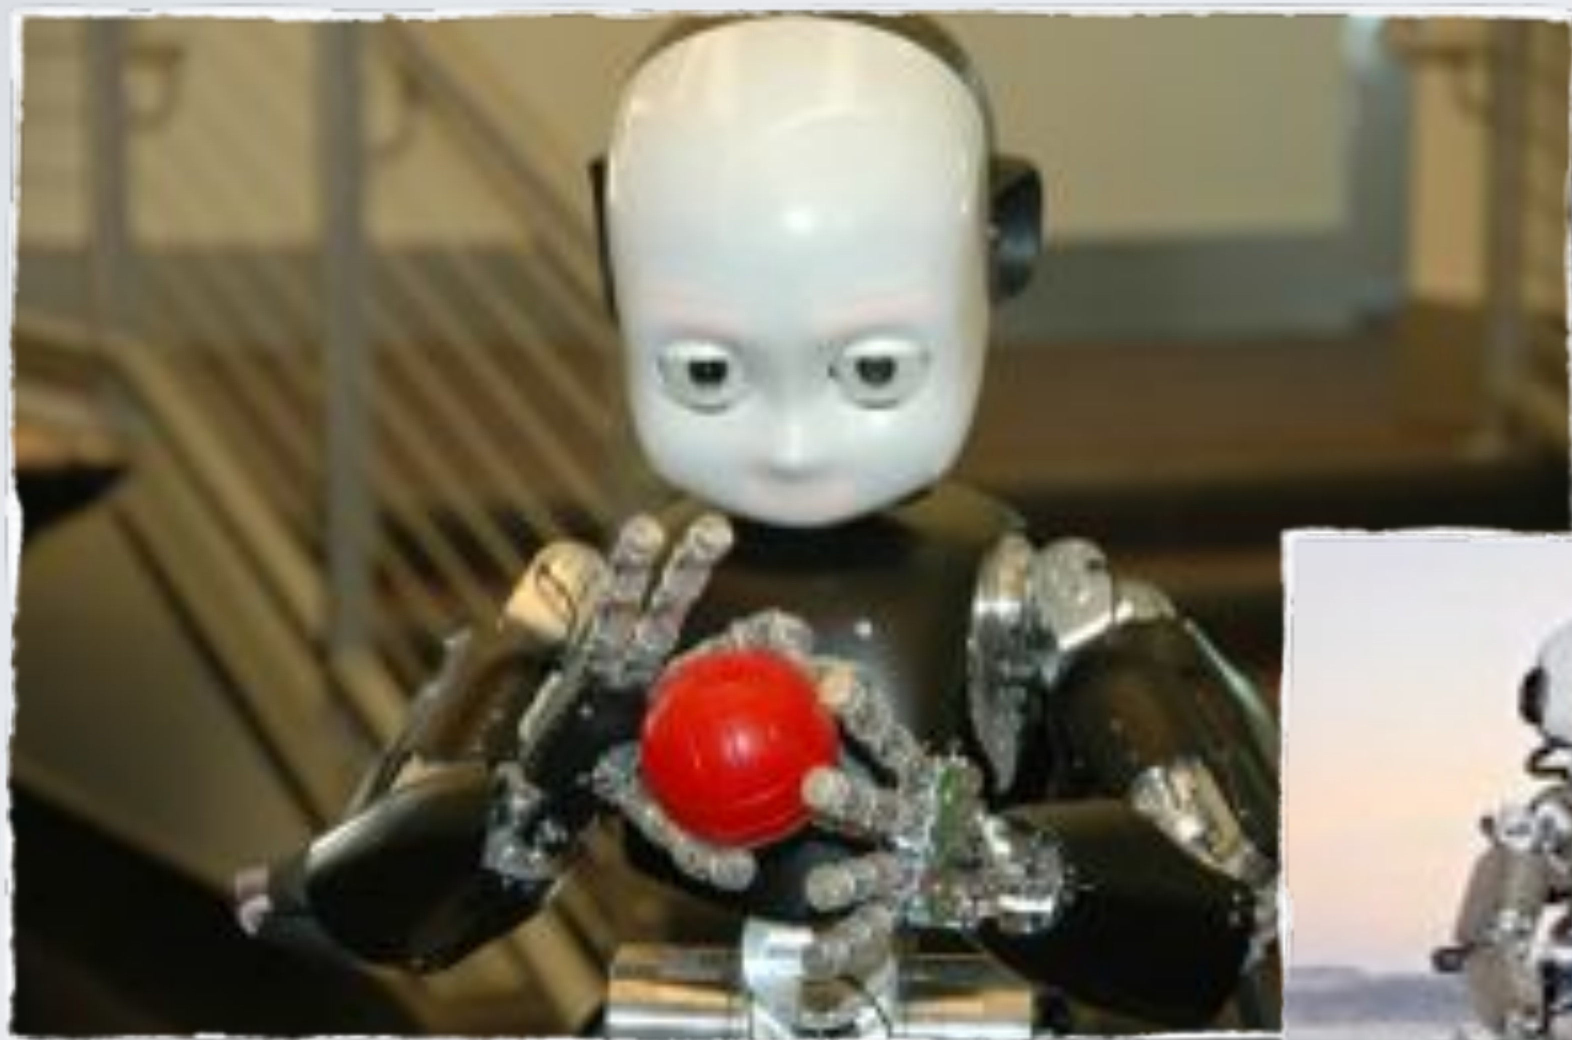

AN APPROACH TO SELF-EXTENSION:

REASONING ABOUT DIFFERENT EPISTEMIC ACTIONS AND UNCERTAINTY FOR AUTONOMOUS KNOWLEDGE GATHERING

Marc Hanheide
University of Lincoln (was Birmingham)

Alper Aydemir
Richard Dearden
Moritz Göbelbecker
Charles Gretton
Nick Hawes
Patric Jensfelt

Andrzej Pronobis
Kristoffer Sjöö
Markus Vincze
Jeremy Wyatt
Michael Zillich
and many more...

Cognitive System that
Self-Understand and Self-Extend

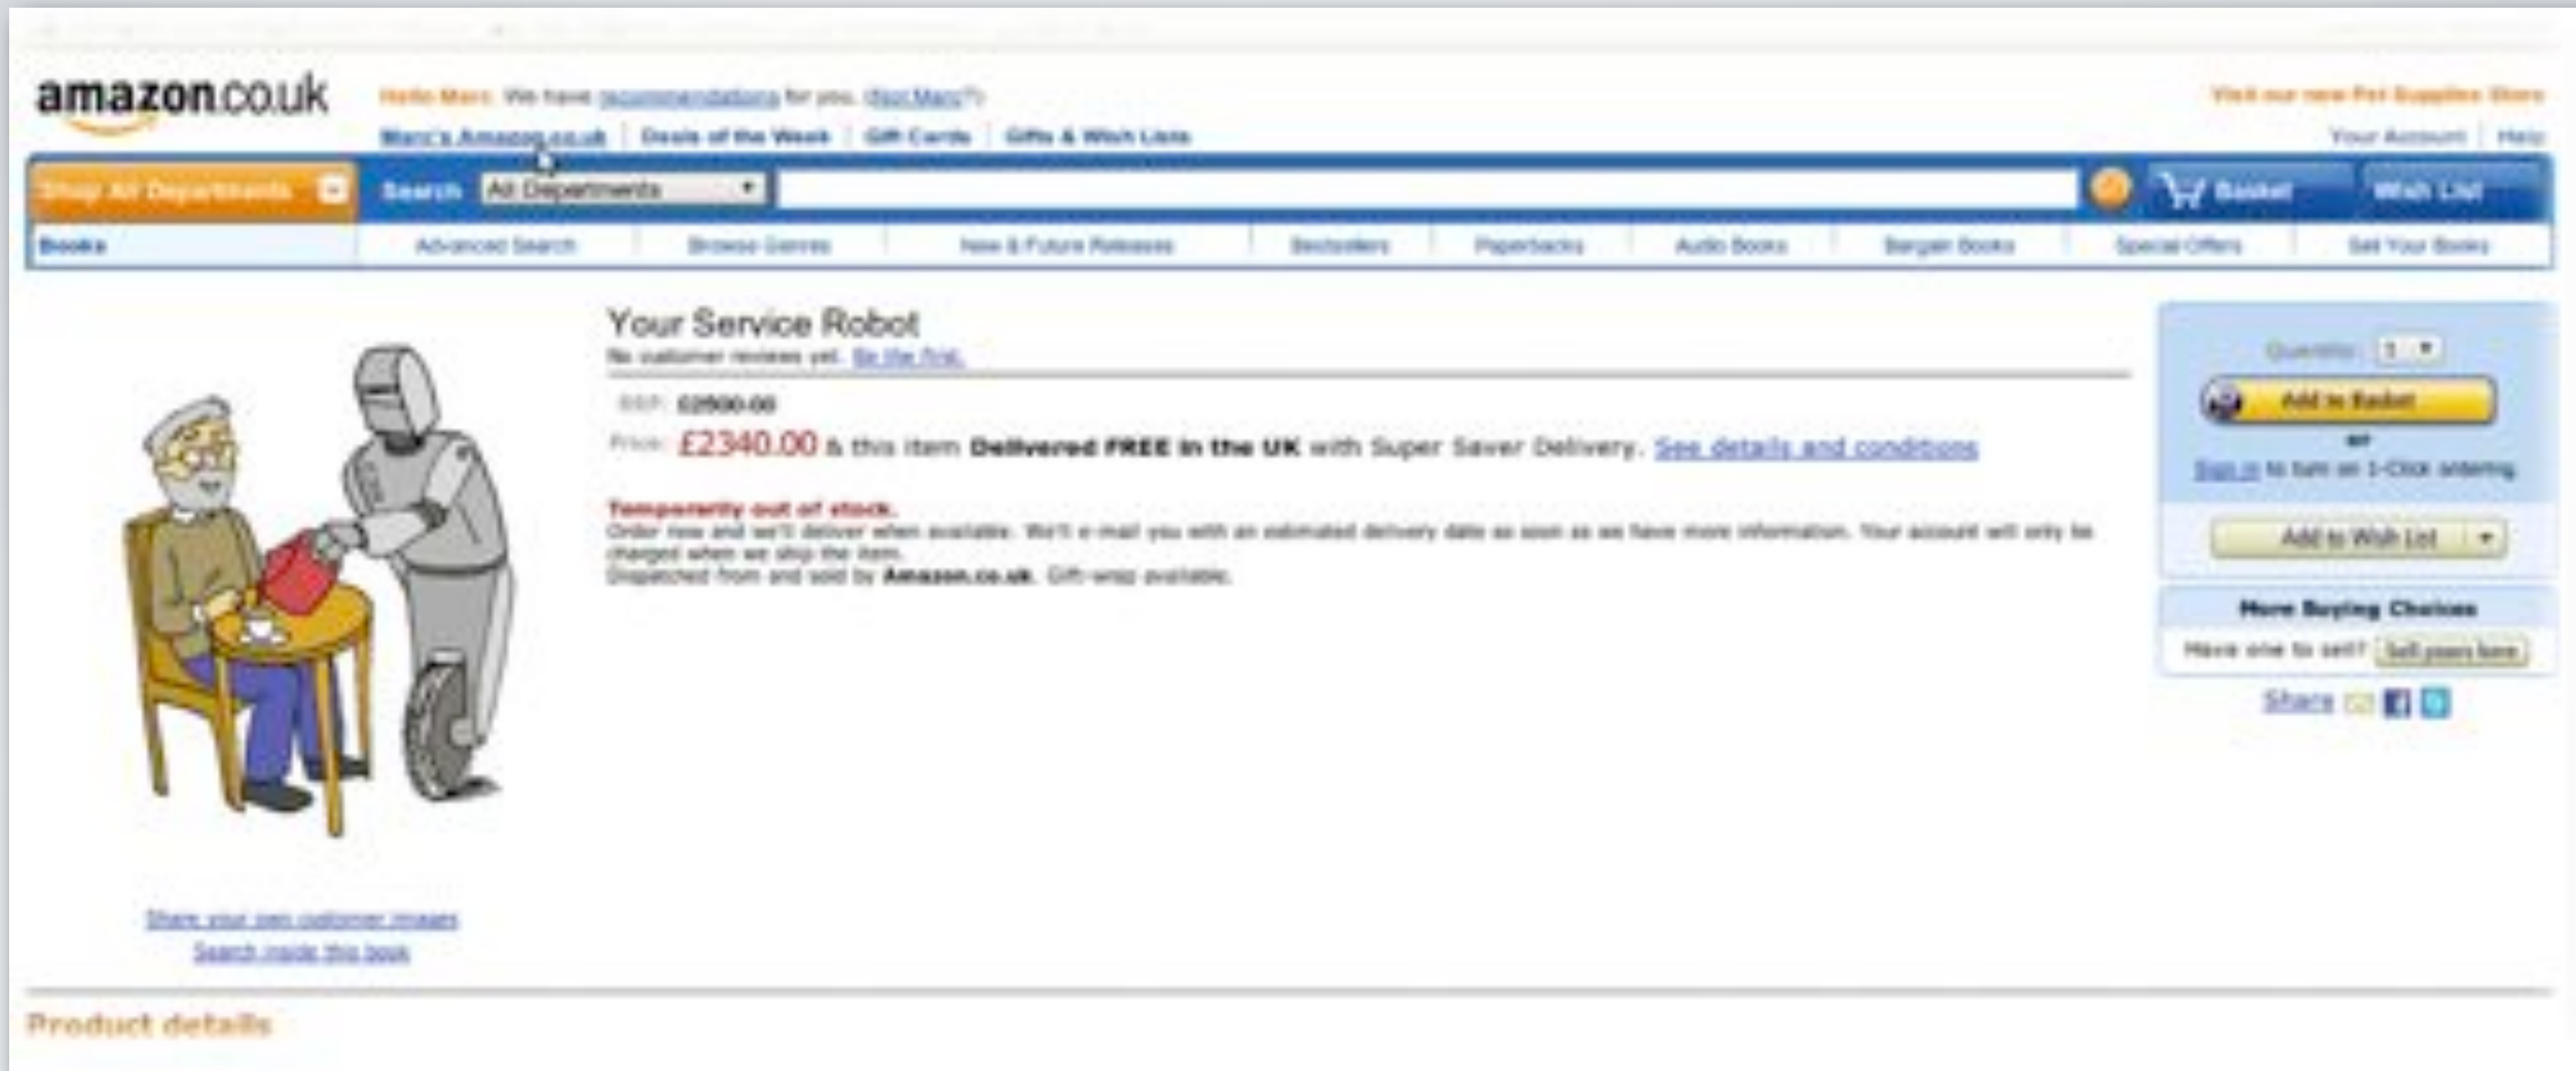
amazon.co.uk
Deals of the Week
Gift Cards
Gifts & Wish Lists
Your Account
Help
Shop All Departments
Search
All Departments
Basket
Wish List
Books
Advanced Search
Browse Genres
New & Future Releases
Bestsellers
Paperbacks
Audio Books
Bargain Books
Special Offers
Sell Your Books
Your Service Robot
£2340.00 & this item Delivered FREE in the UK with Super Saver Delivery. See details and conditions
Temporarily out of stock.
Dispatched from and sold by Amazon.co.uk. Gift-wrap available.
Add to Basket
Add to Wish List
More Buying Choices
Share
Product details

ABILITIES

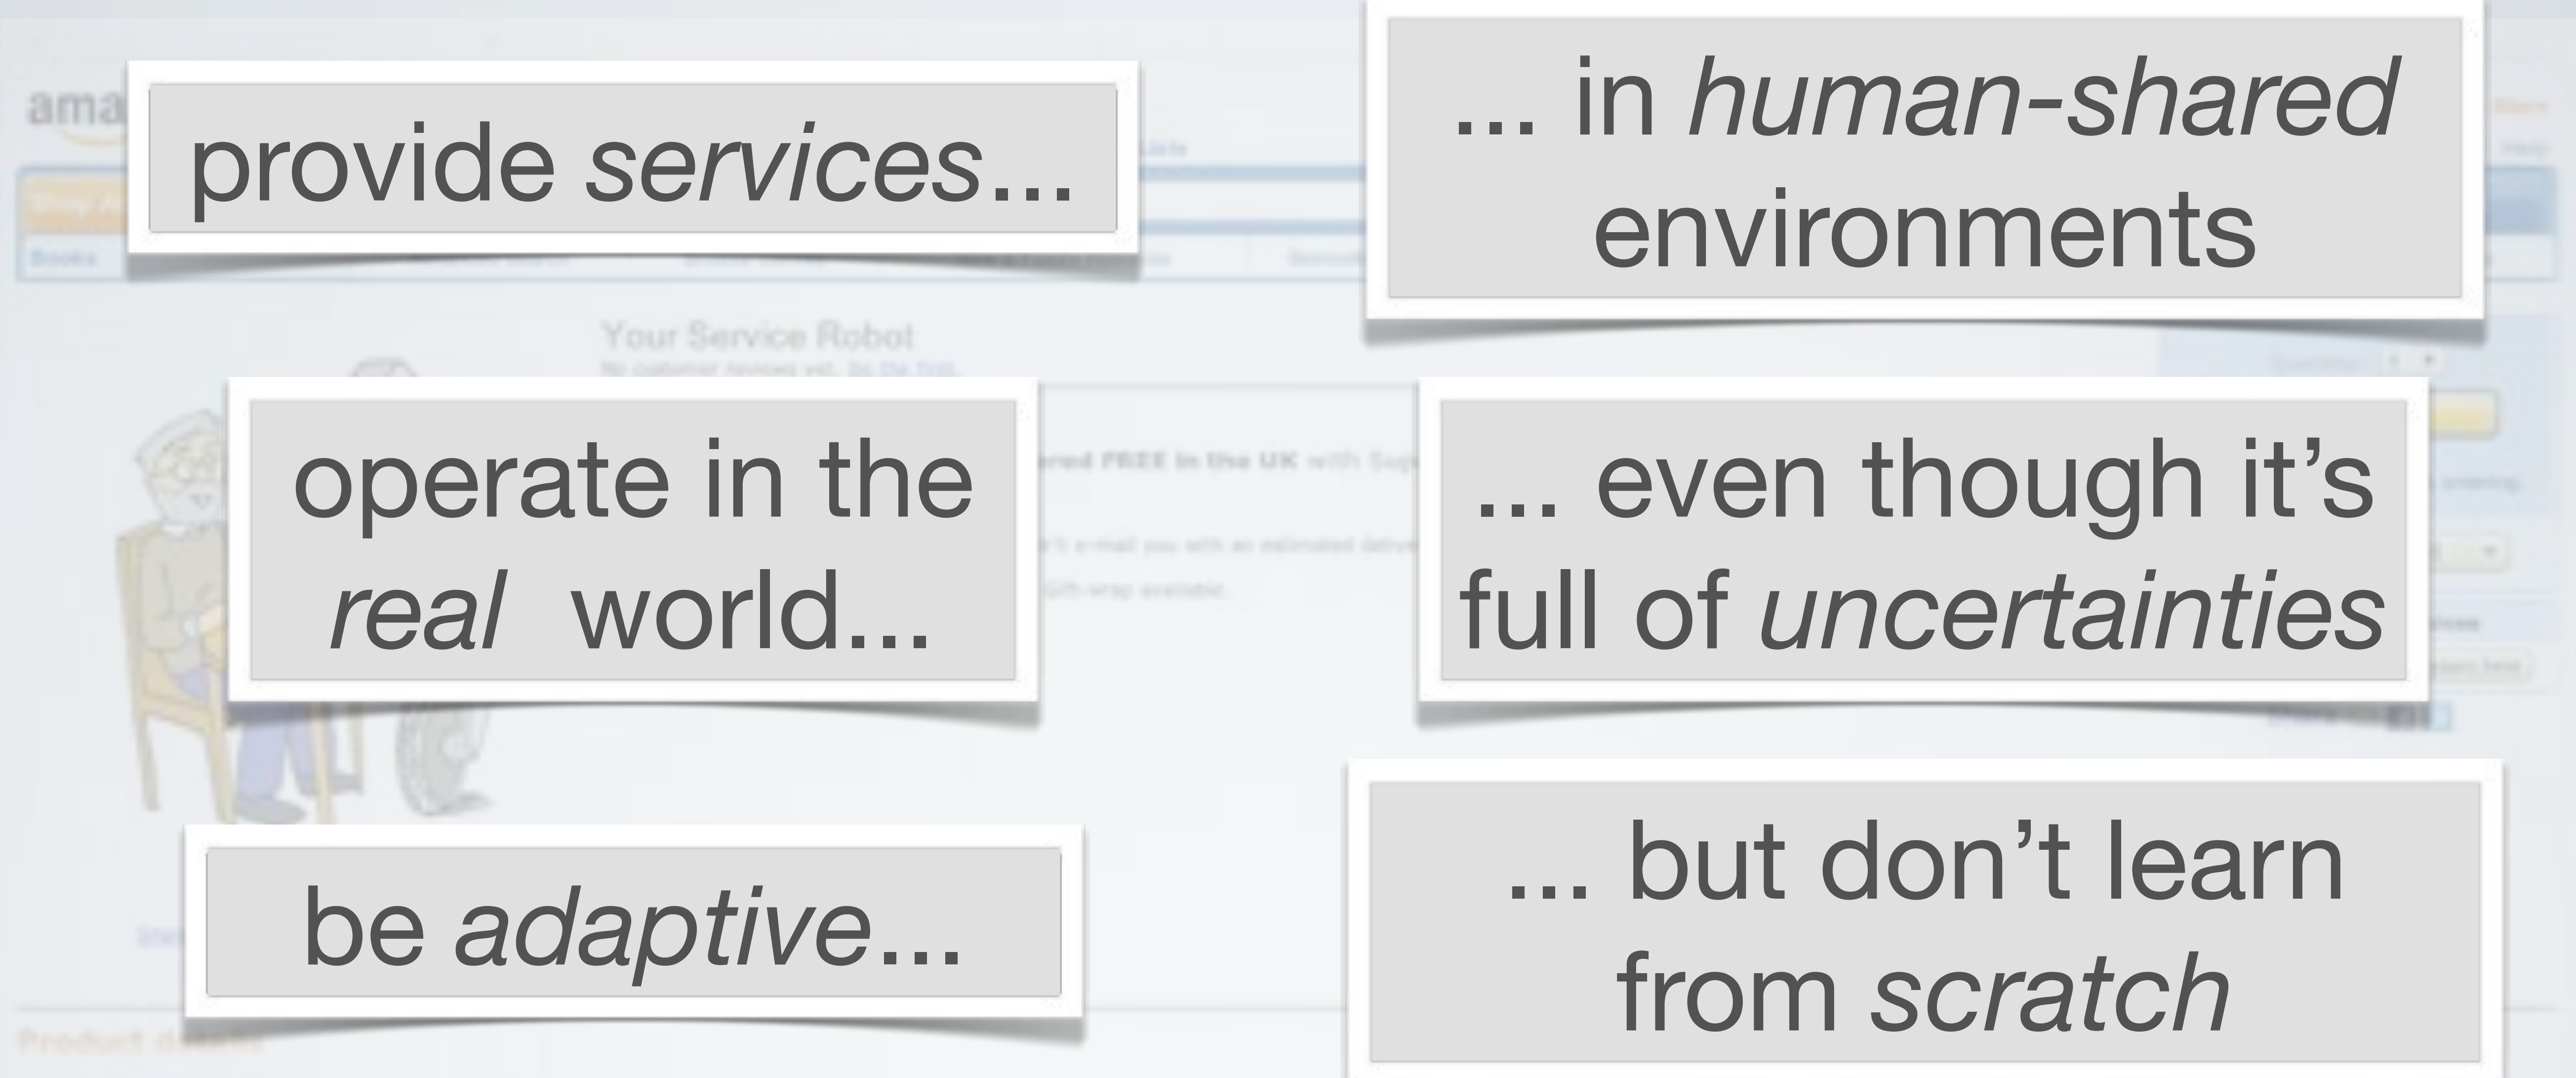

provide *services*...

... in *human-shared* environments

operate in the *real* world...

... even though it's full of *uncertainties*

be *adaptive*...

... but don't learn from *scratch*

ARE YOU HUNGRY?

> 45000 views

Kunze, L. et al., 2012. Searching Objects in Large-scale Indoor Environments: A Decision-theoretic Approach. In IEEE International Conference on Robotics and Automation. pp. 4385-4390.

ARE YOU HUNGRY?

> 45000 views

Kunze, L. et al., 2012. Searching Objects in Large-scale Indoor Environments: A Decision-theoretic Approach. In IEEE International Conference on Robotics and Automation. pp. 4385-4390.

WHAT IF...

... the object is not at (X,Y,Z)?

... the robot wouldn't know where the shop/ fridge is?

... the map would be unknown (no global coordinates)?

... sensing would be uncertain?

... the common-sense knowledge would be wrong?

Kunze, L. et al., 2012. Searching Objects in Large-scale Indoor Environments: A Decision-theoretic Approach. In IEEE International Conference on Robotics and Automation. pp. 4385-4390.

A ROBOT IN MY (OLD) HOME

... th

... s

bo

ge

?

“Essentially a T-shirt on a stick with a wheeled base and camera for a face, Dora’s *draw is her brain*.”

[New Scientist]

IF... THEN!

Autonomous Visual Search

e object is not at (X,Y,Z)?

... the map would be unknown (no global coordinates)?

Autonomous Exploration

... sensing would be uncertain?

Probabilistic Planning

Room Categorisation

... the robot wouldn't know where the shop/ kitchen is?

... the common-sense knowledge would be wrong?

Interactive Hypothesis Verification

Hanheide, M. et al., 2011. Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour. In Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI).

DORA SEARCHING

Autonomous Visual Search

HOW DOES DORA SEE OBJECTS?

- 2D image SIFTs + 3D surface SIFTs correspondence
- RANSAC based pose estimation [Prankl, 2010]
- based on *known* CAD model and *texture*
- requires pre-trained models

PLANNING TO SEARCH

- Take into account the characteristics of CV operators
- Tsotsos et al. (2010)
- Limitations:
  - range
  - pre-mapped
  - focus

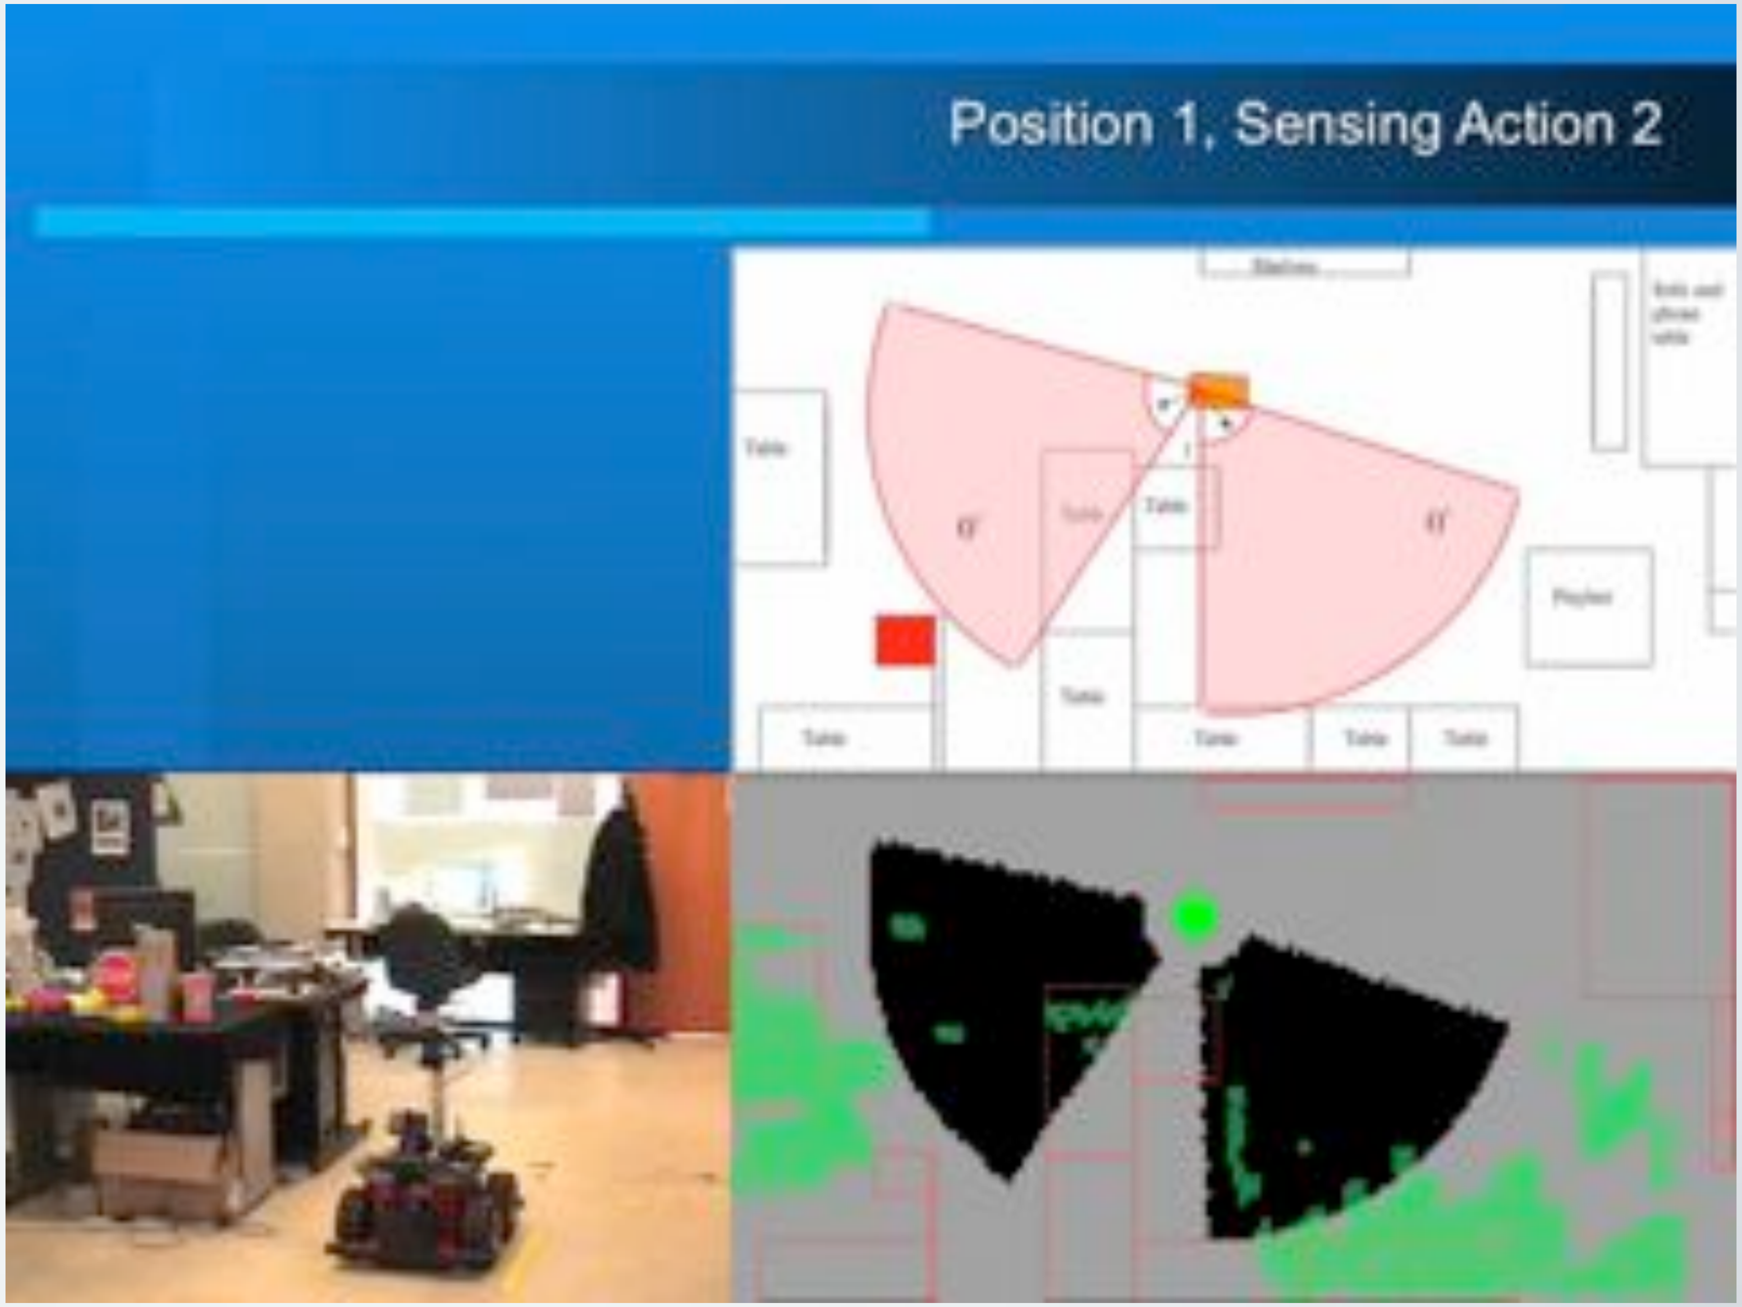


John's cup
office stapler
•cornflakes
•bowl
•milk
mail
1984 Book
Mary's laptop

MORE KNOWLEDGE

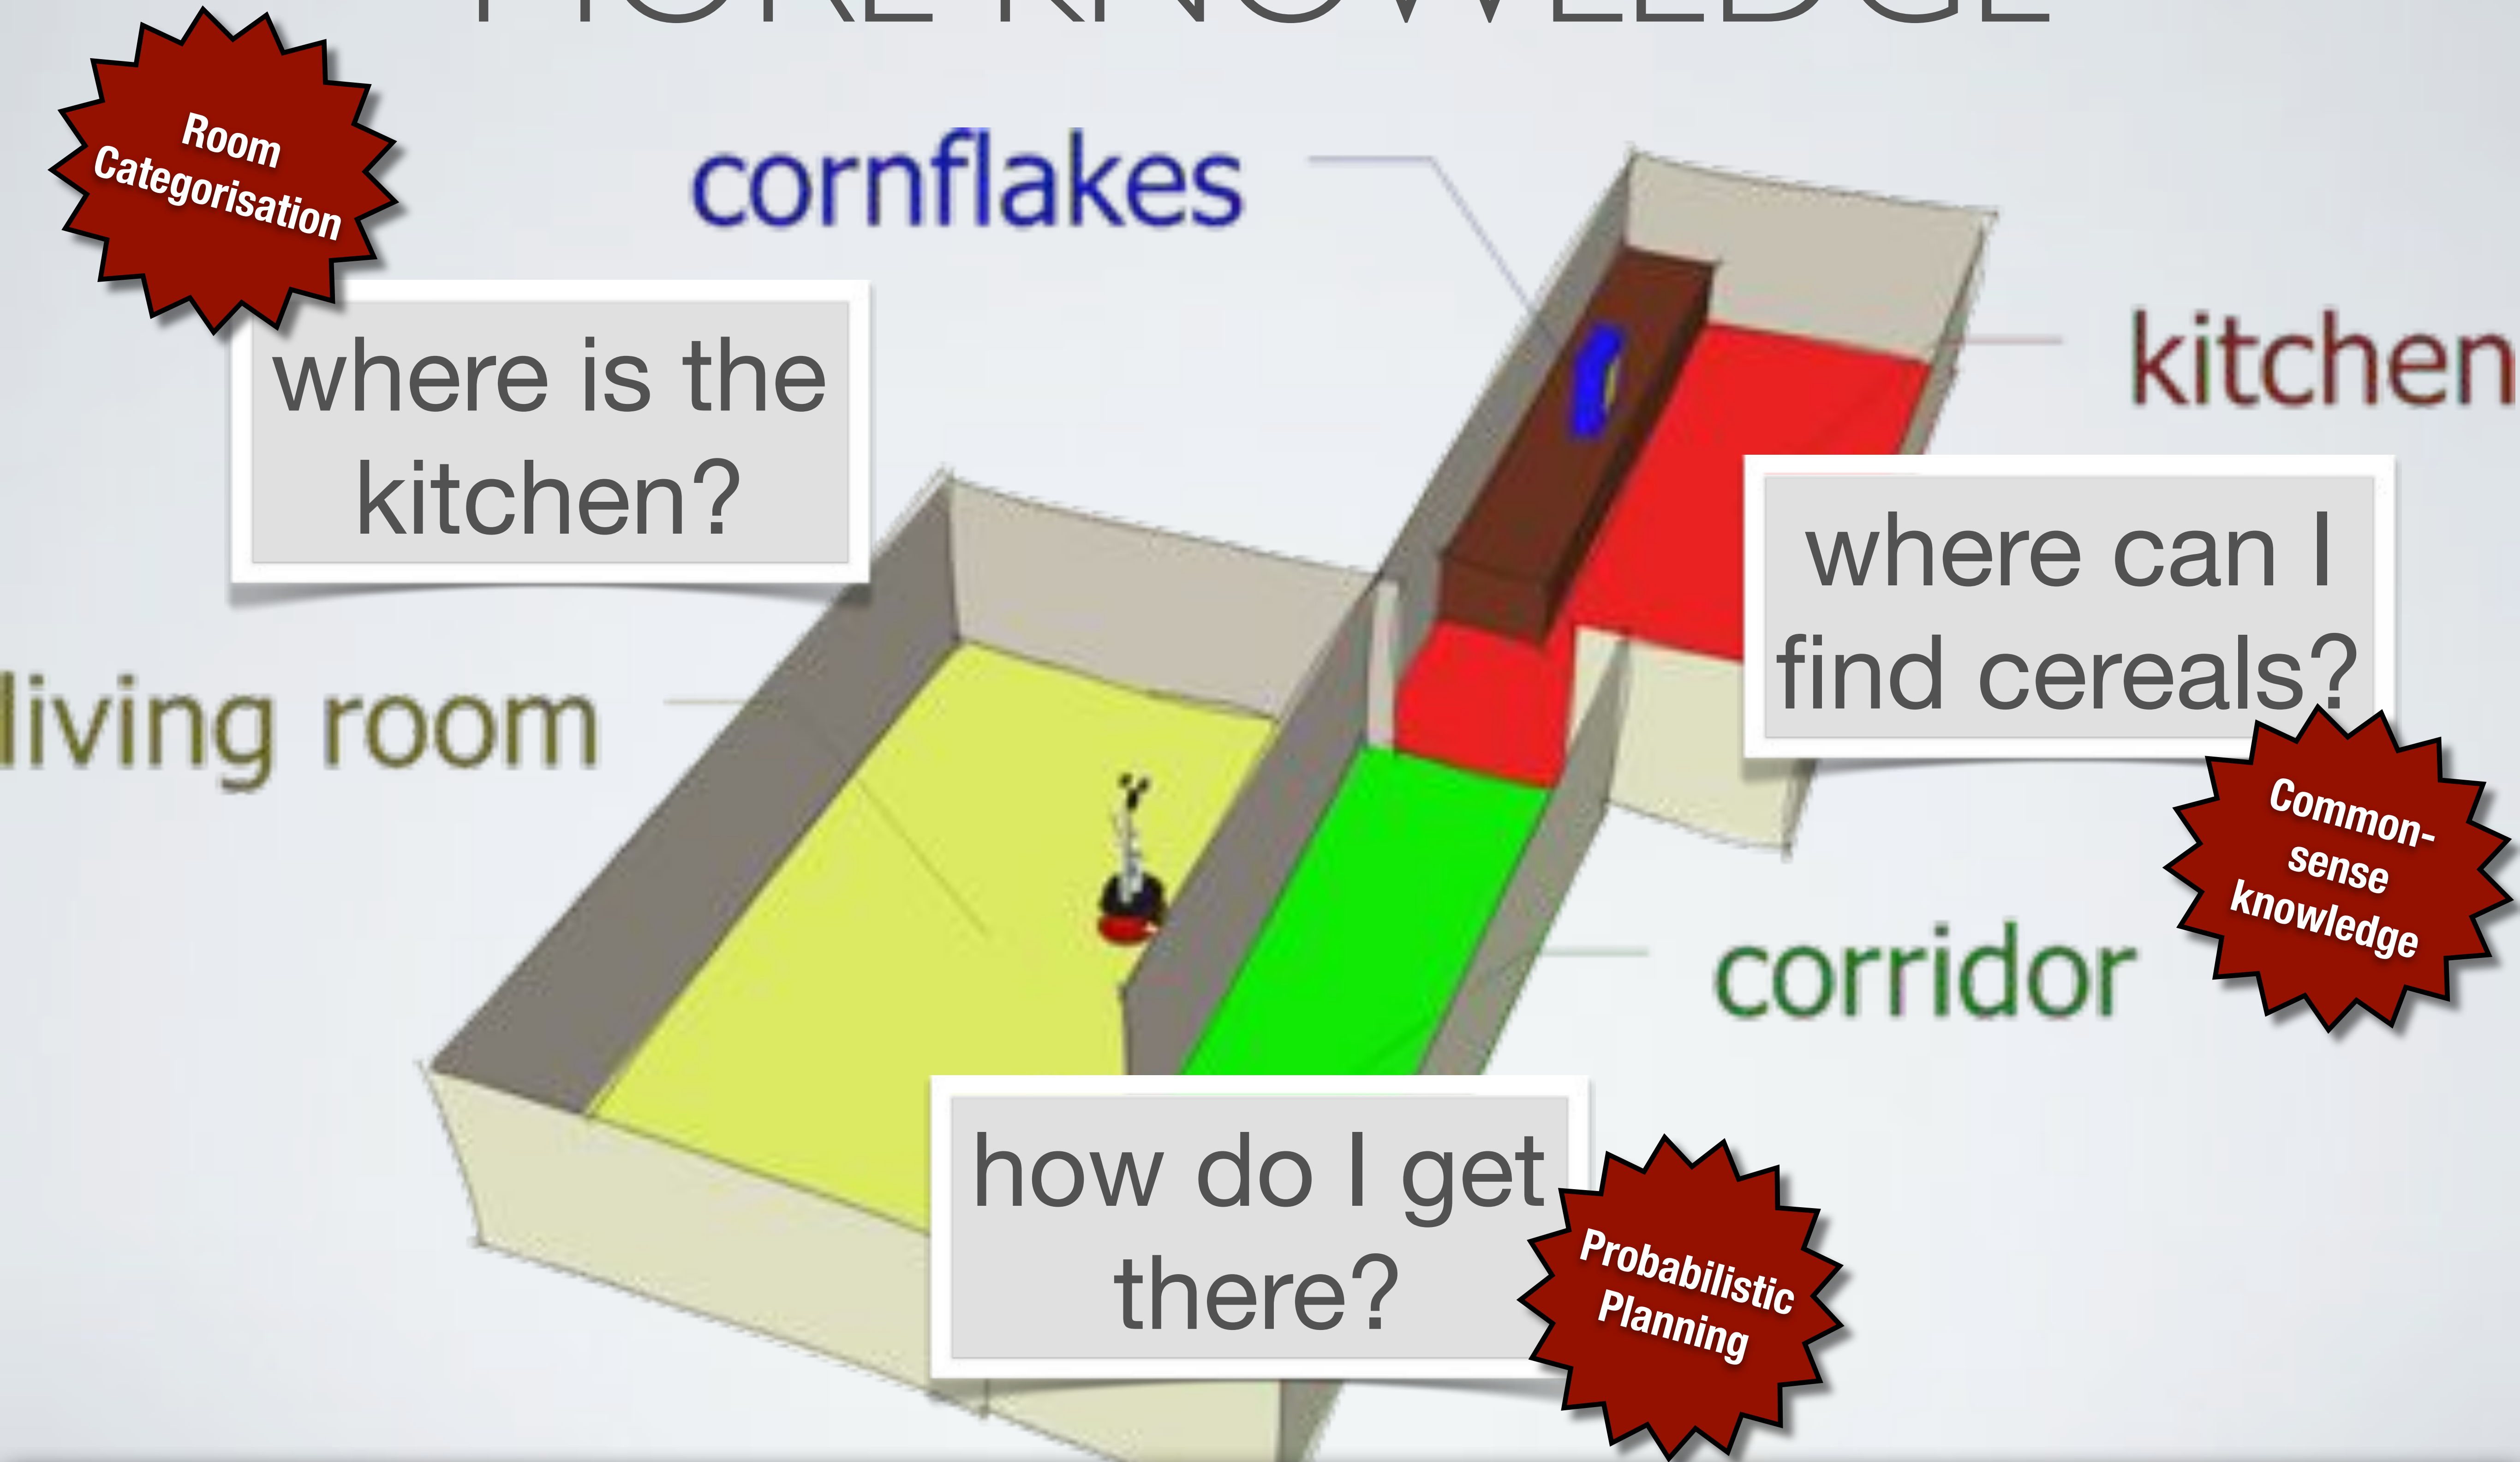

Hanheide, M. et al. Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour. *IJCAI 2011*.

WHERE TO SEARCH?

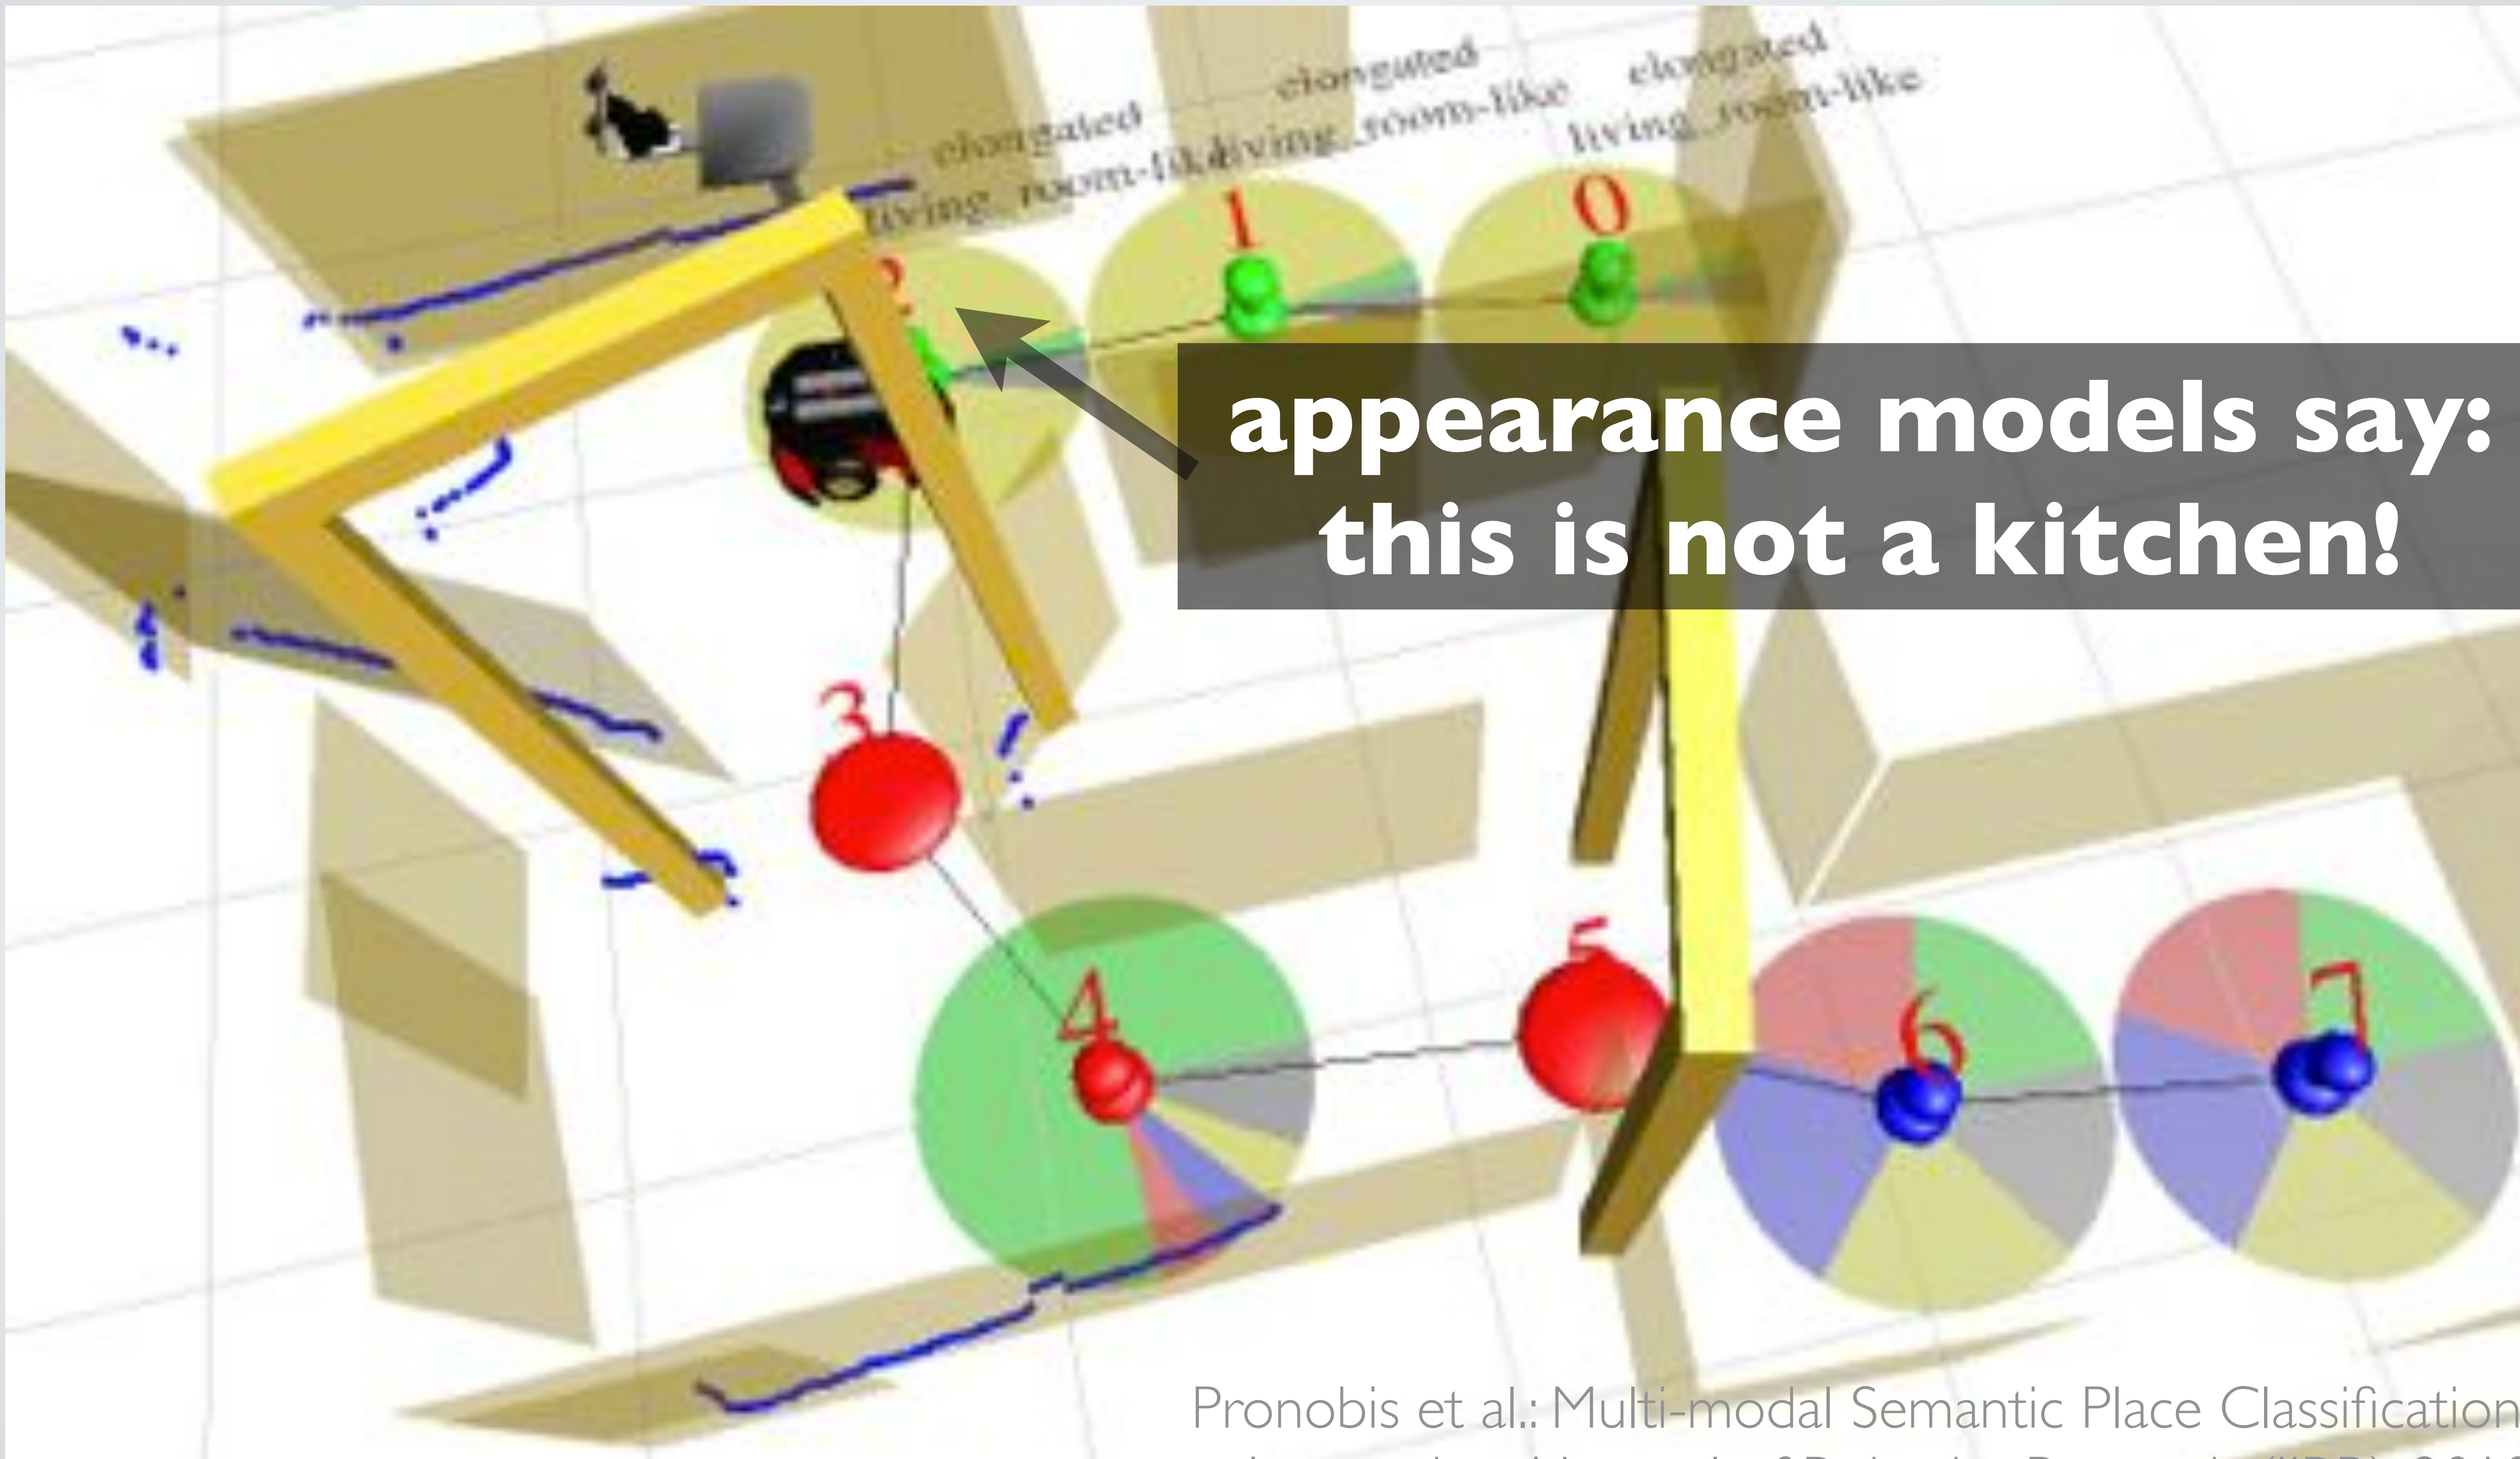

appearance models say: this is not a kitchen!

Pronobis et al.: Multi-modal Semantic Place Classification, International Journal of Robotics Research (IJRR), 2010

DORA

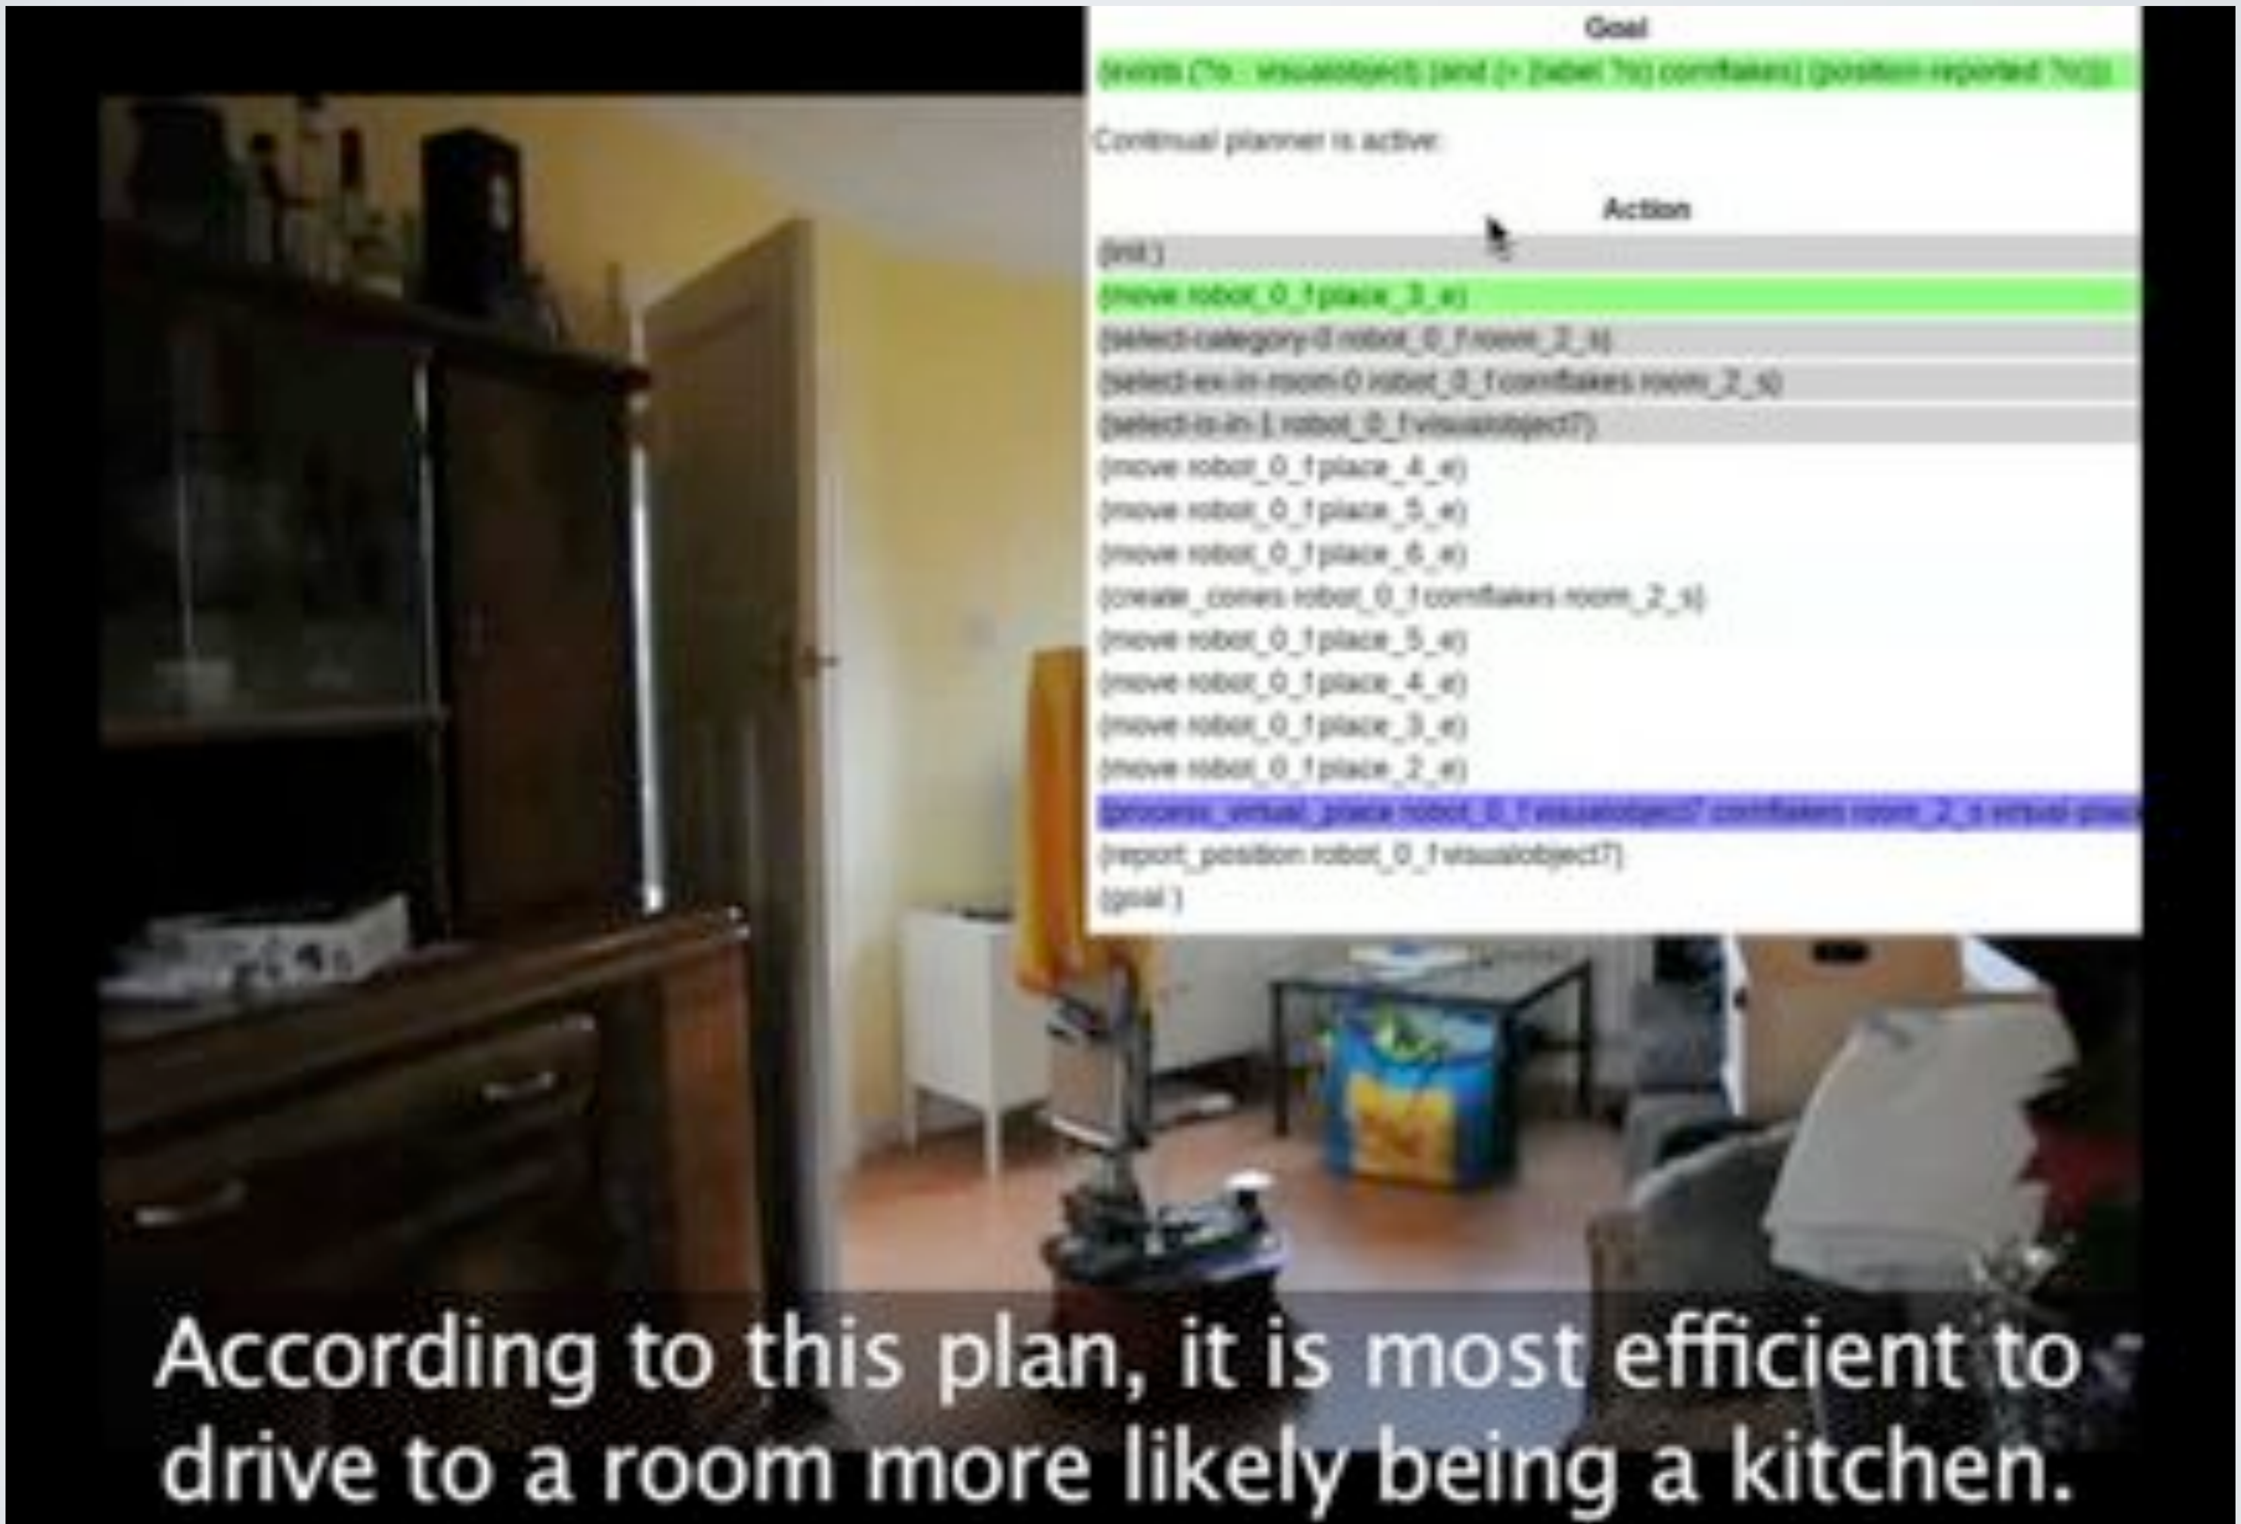
Goal
Continual planner is active
Action
According to this plan, it is most efficient to drive to a room more likely being a kitchen.

DORA SEARCHES THE KITCHEN

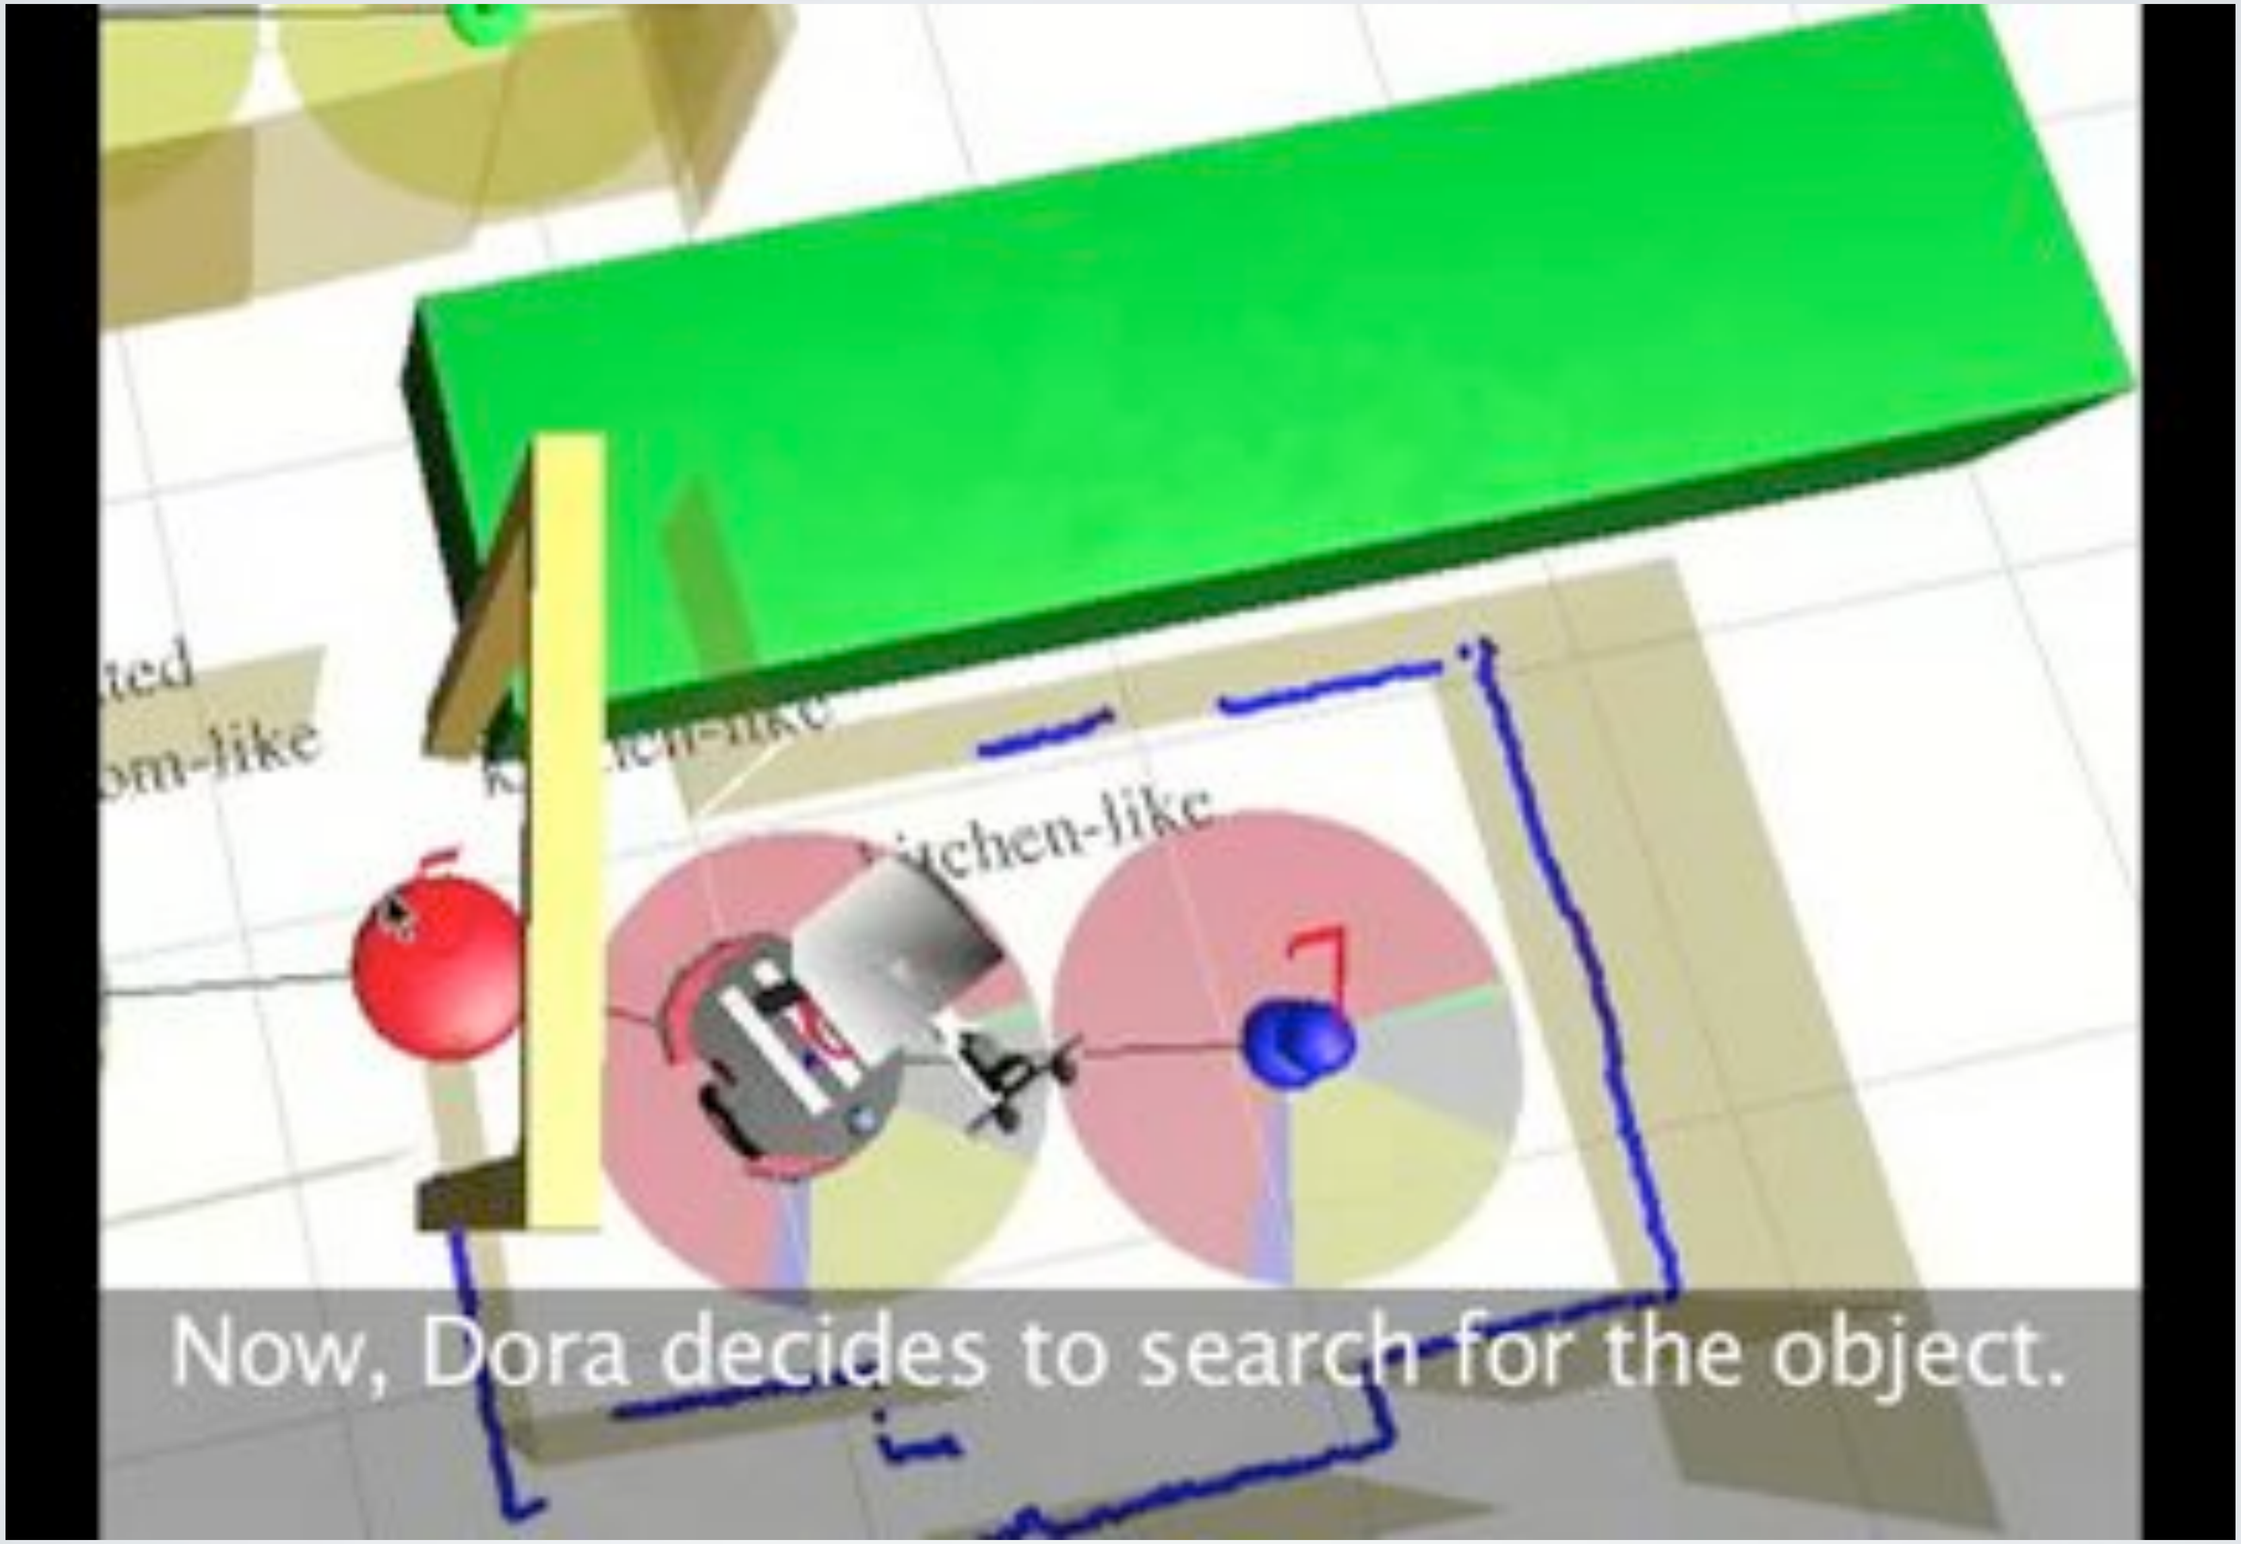

RESULTS

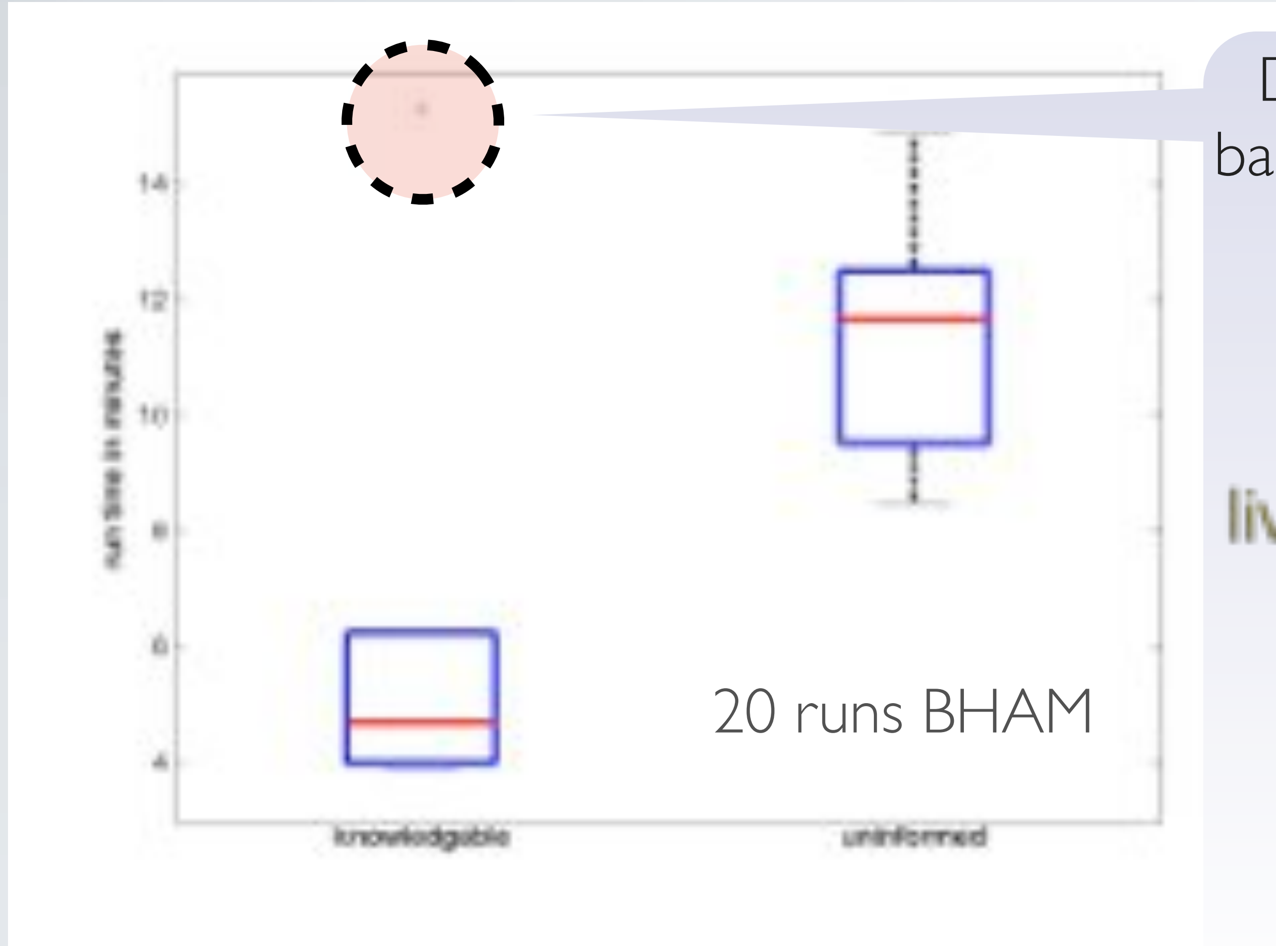

Dora failed to find the object in first place, went back to the living room and returned to the kitchen

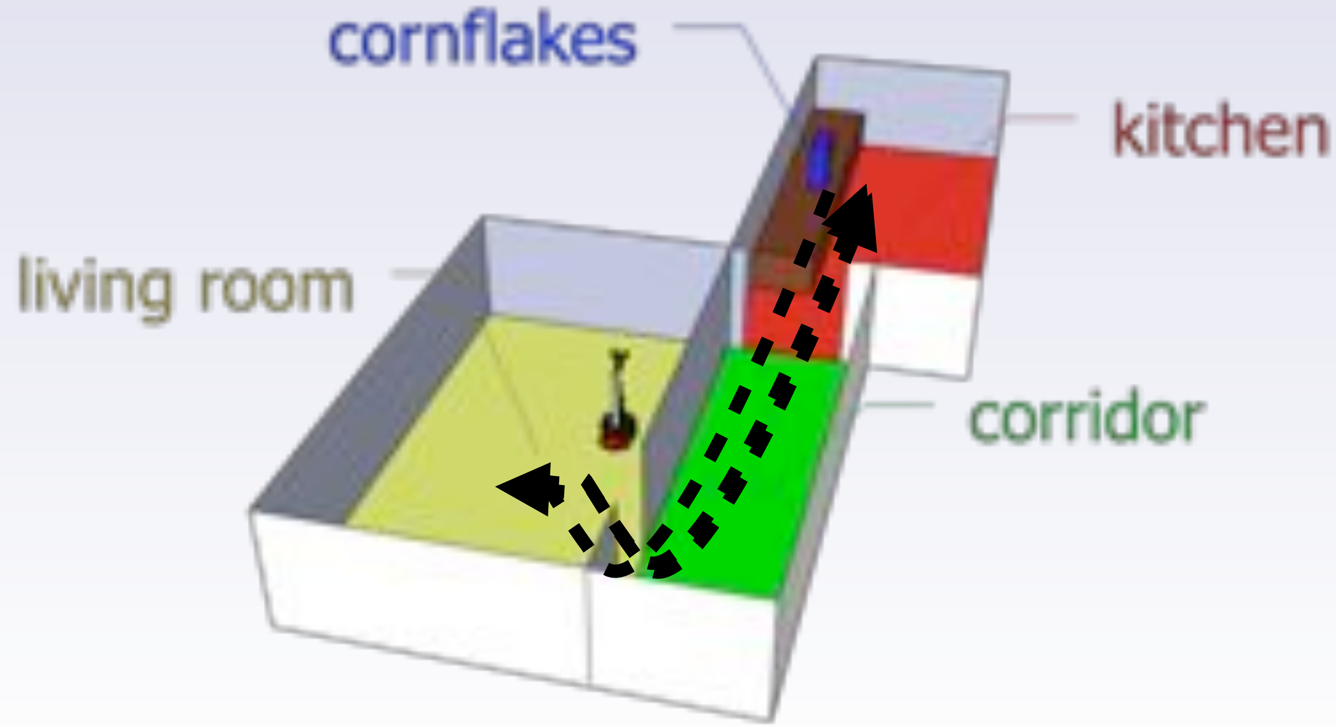

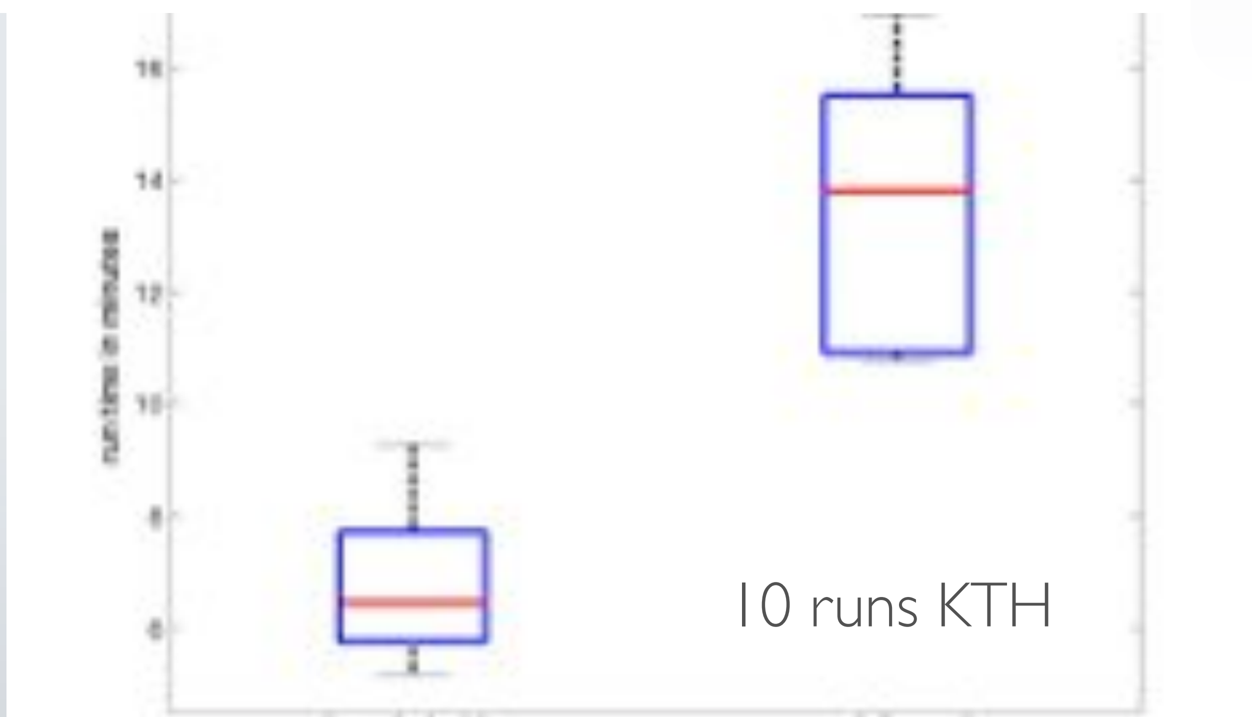

... sensing would be uncertain?

THE MESSY STUDENT APARTMENT

... the common-sense knowledge would be wrong?

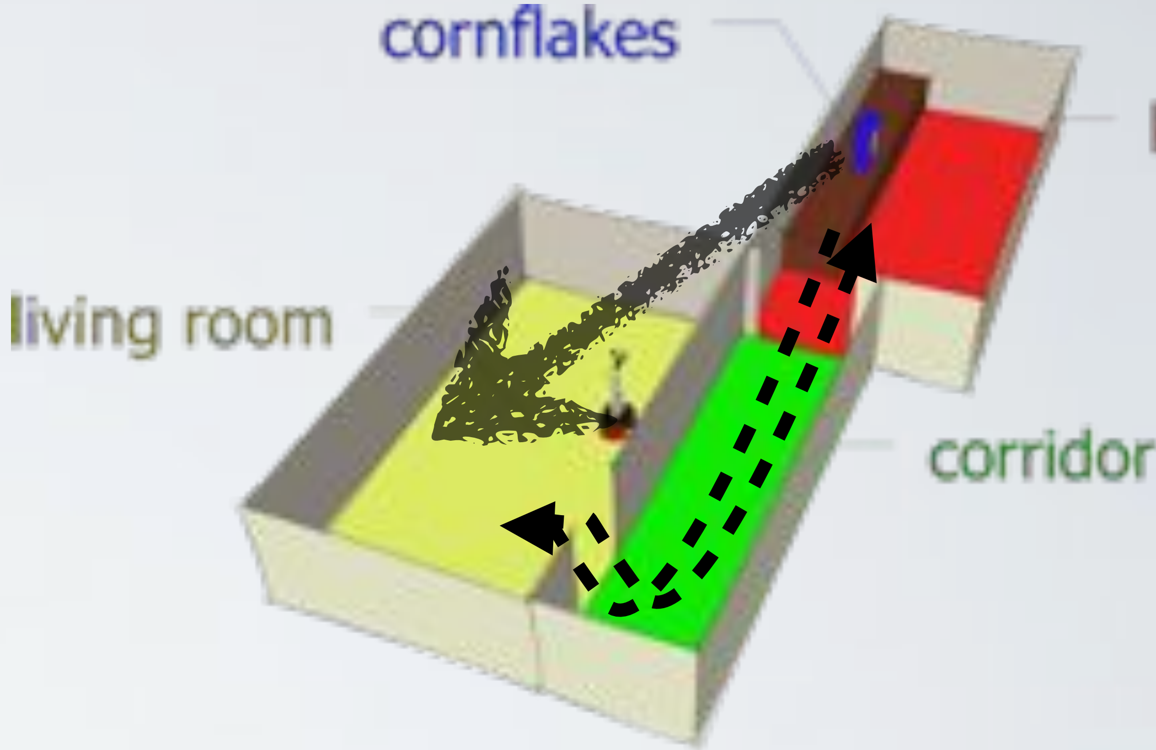

conf.	obj. loc.	lesion	♯succ./♯tot. *H*	avg. time *H*
FC	kitchen	no	10/10	5.8min
LC	kitchen	yes	9/10	11.3min
FNC	liv. room,	no	3/3	10.2min

DEALING WITH UNCERTAINTY

- sensing algorithms can fail
  - not seeing the object
  - mis-classifying a room
- common-sense knowledge can be wrong, even though useful in most cases

You know that, do you?

it’s all so uncertain... how to deal with it?

Hanheide et al.: IJCAI, 2011

THE COGX APPROACH

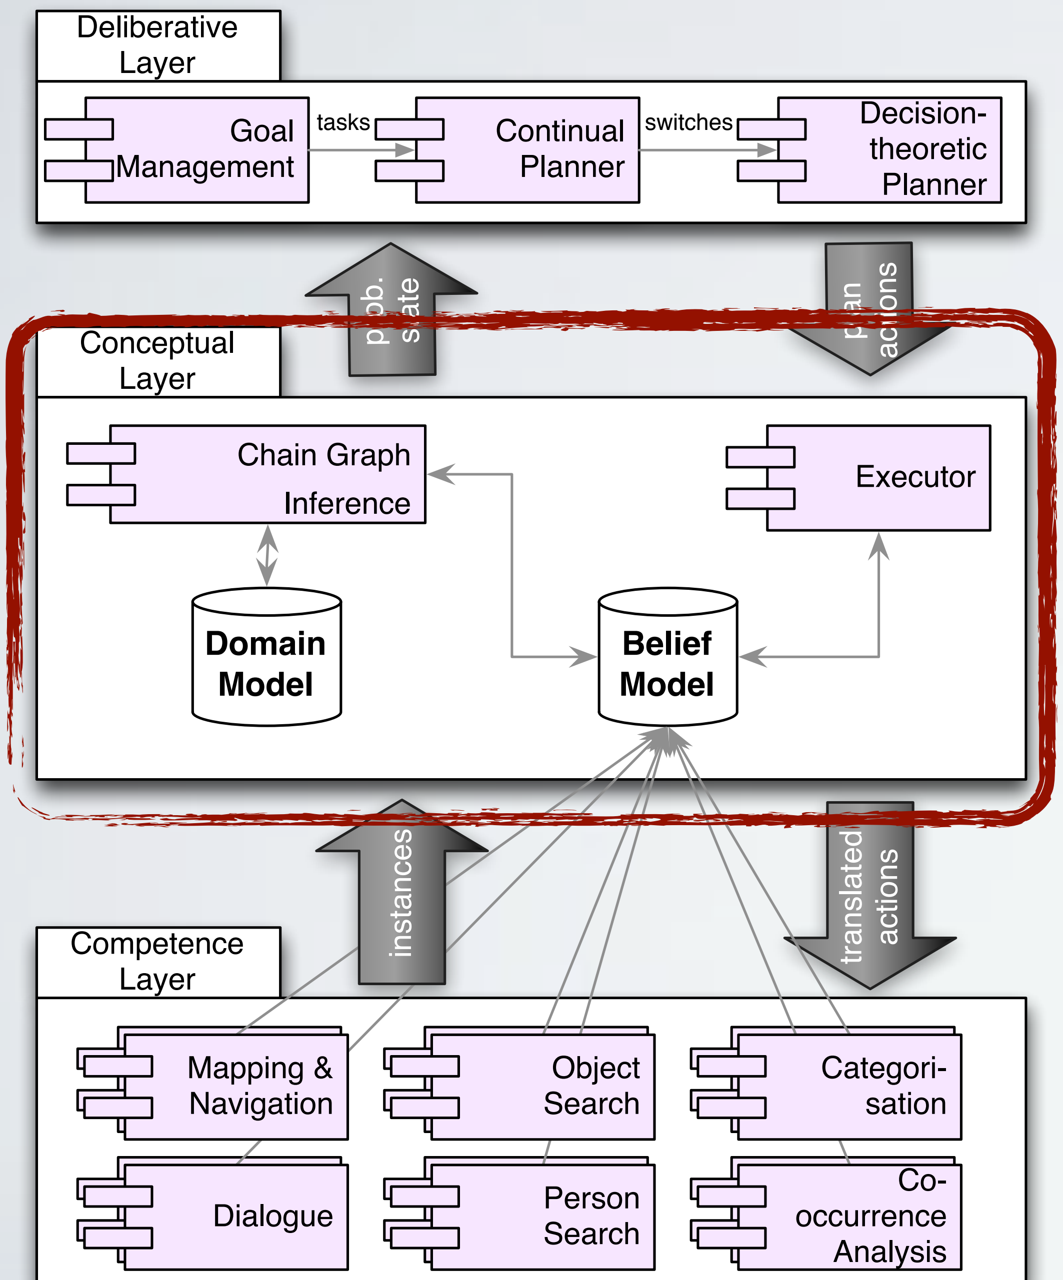

- domain-independent deliberative layer
  - plan autonomous behaviour to accomplish various goals in complex, probabilistic state spaces
- domain-dependent conceptual layer
  - accommodates probabilistic domain and instance knowledge
- competence layer
  - implements sensing and acting
- implemented using CAST (not YARP)

http://www.cs.bham.ac.uk/research/projects/cosy/cast/

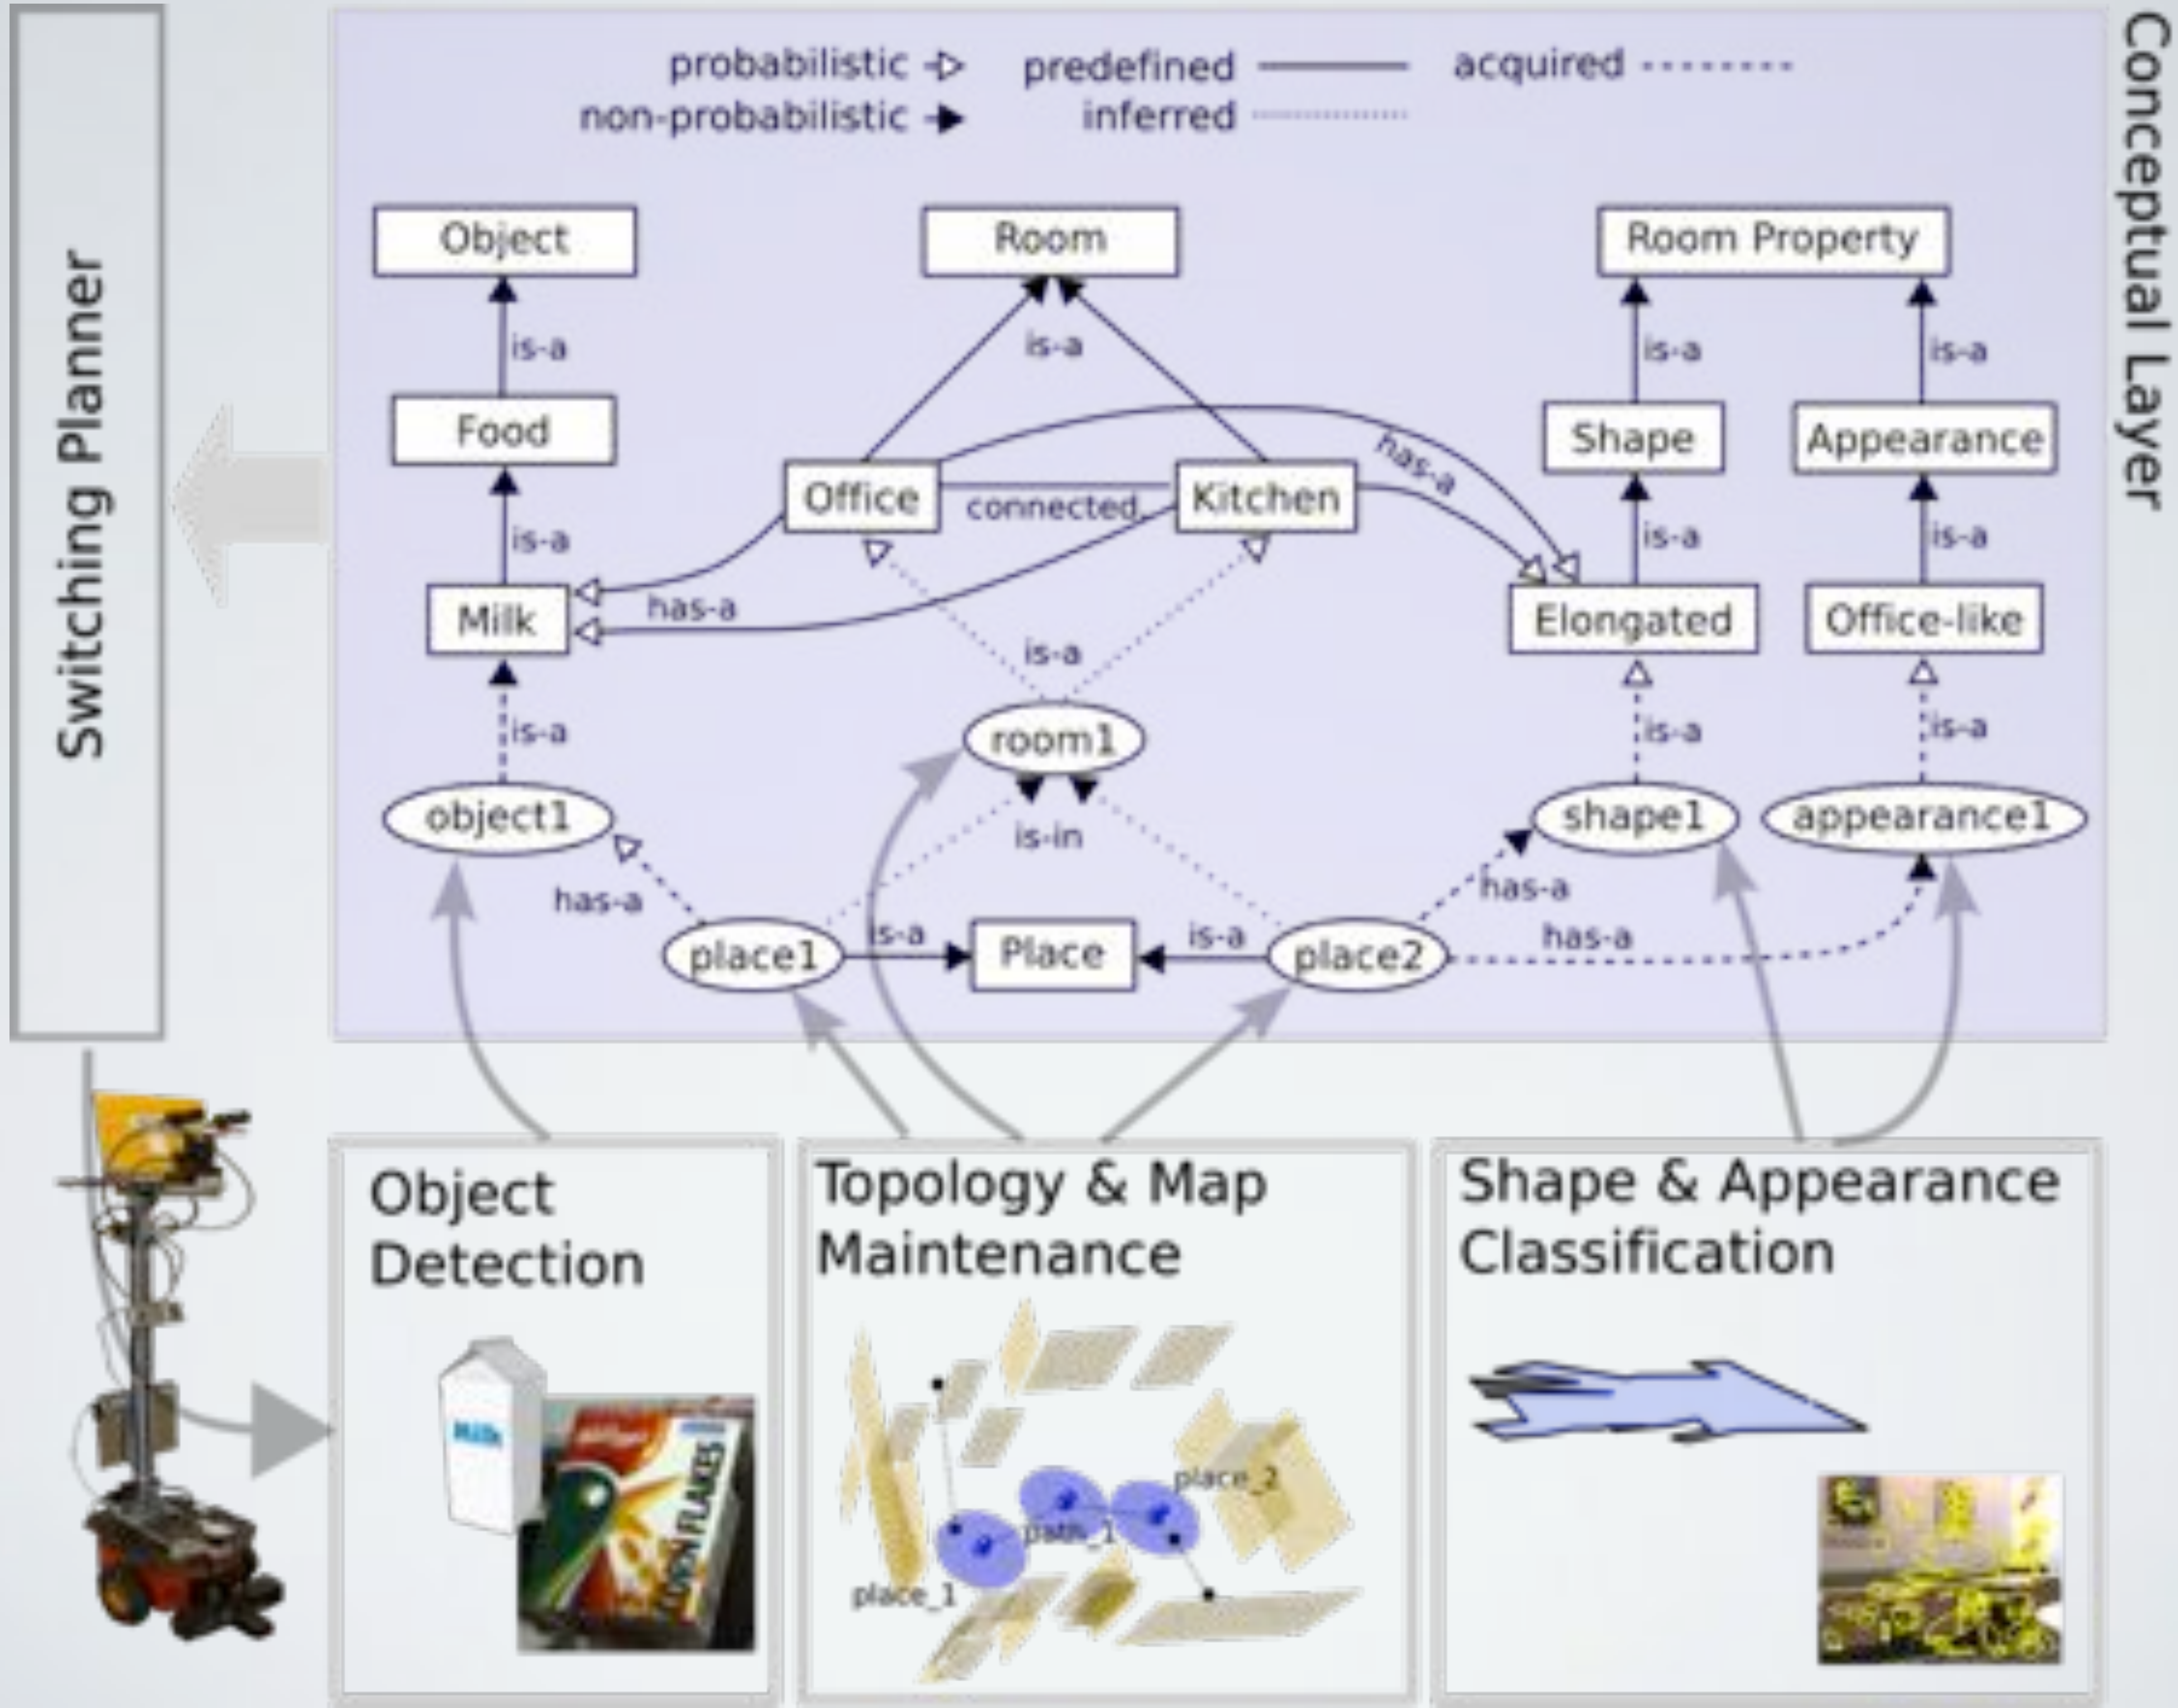

Switching Planner
Conceptual Layer
probabilistic
non-probabilistic
predefined
inferred
acquired
Object
Room
Room Property
is-a
Food
Office
connected
Kitchen
Shape
Appearance
has-a
Milk
Elongated
Office-like
room1
object1
is-in
shape1
appearance1
place1
Place
place2
Object Detection
Topology & Map Maintenance
Shape & Appearance Classification
place_2
path_1
place_1

QUANTIFYING RELATIONS

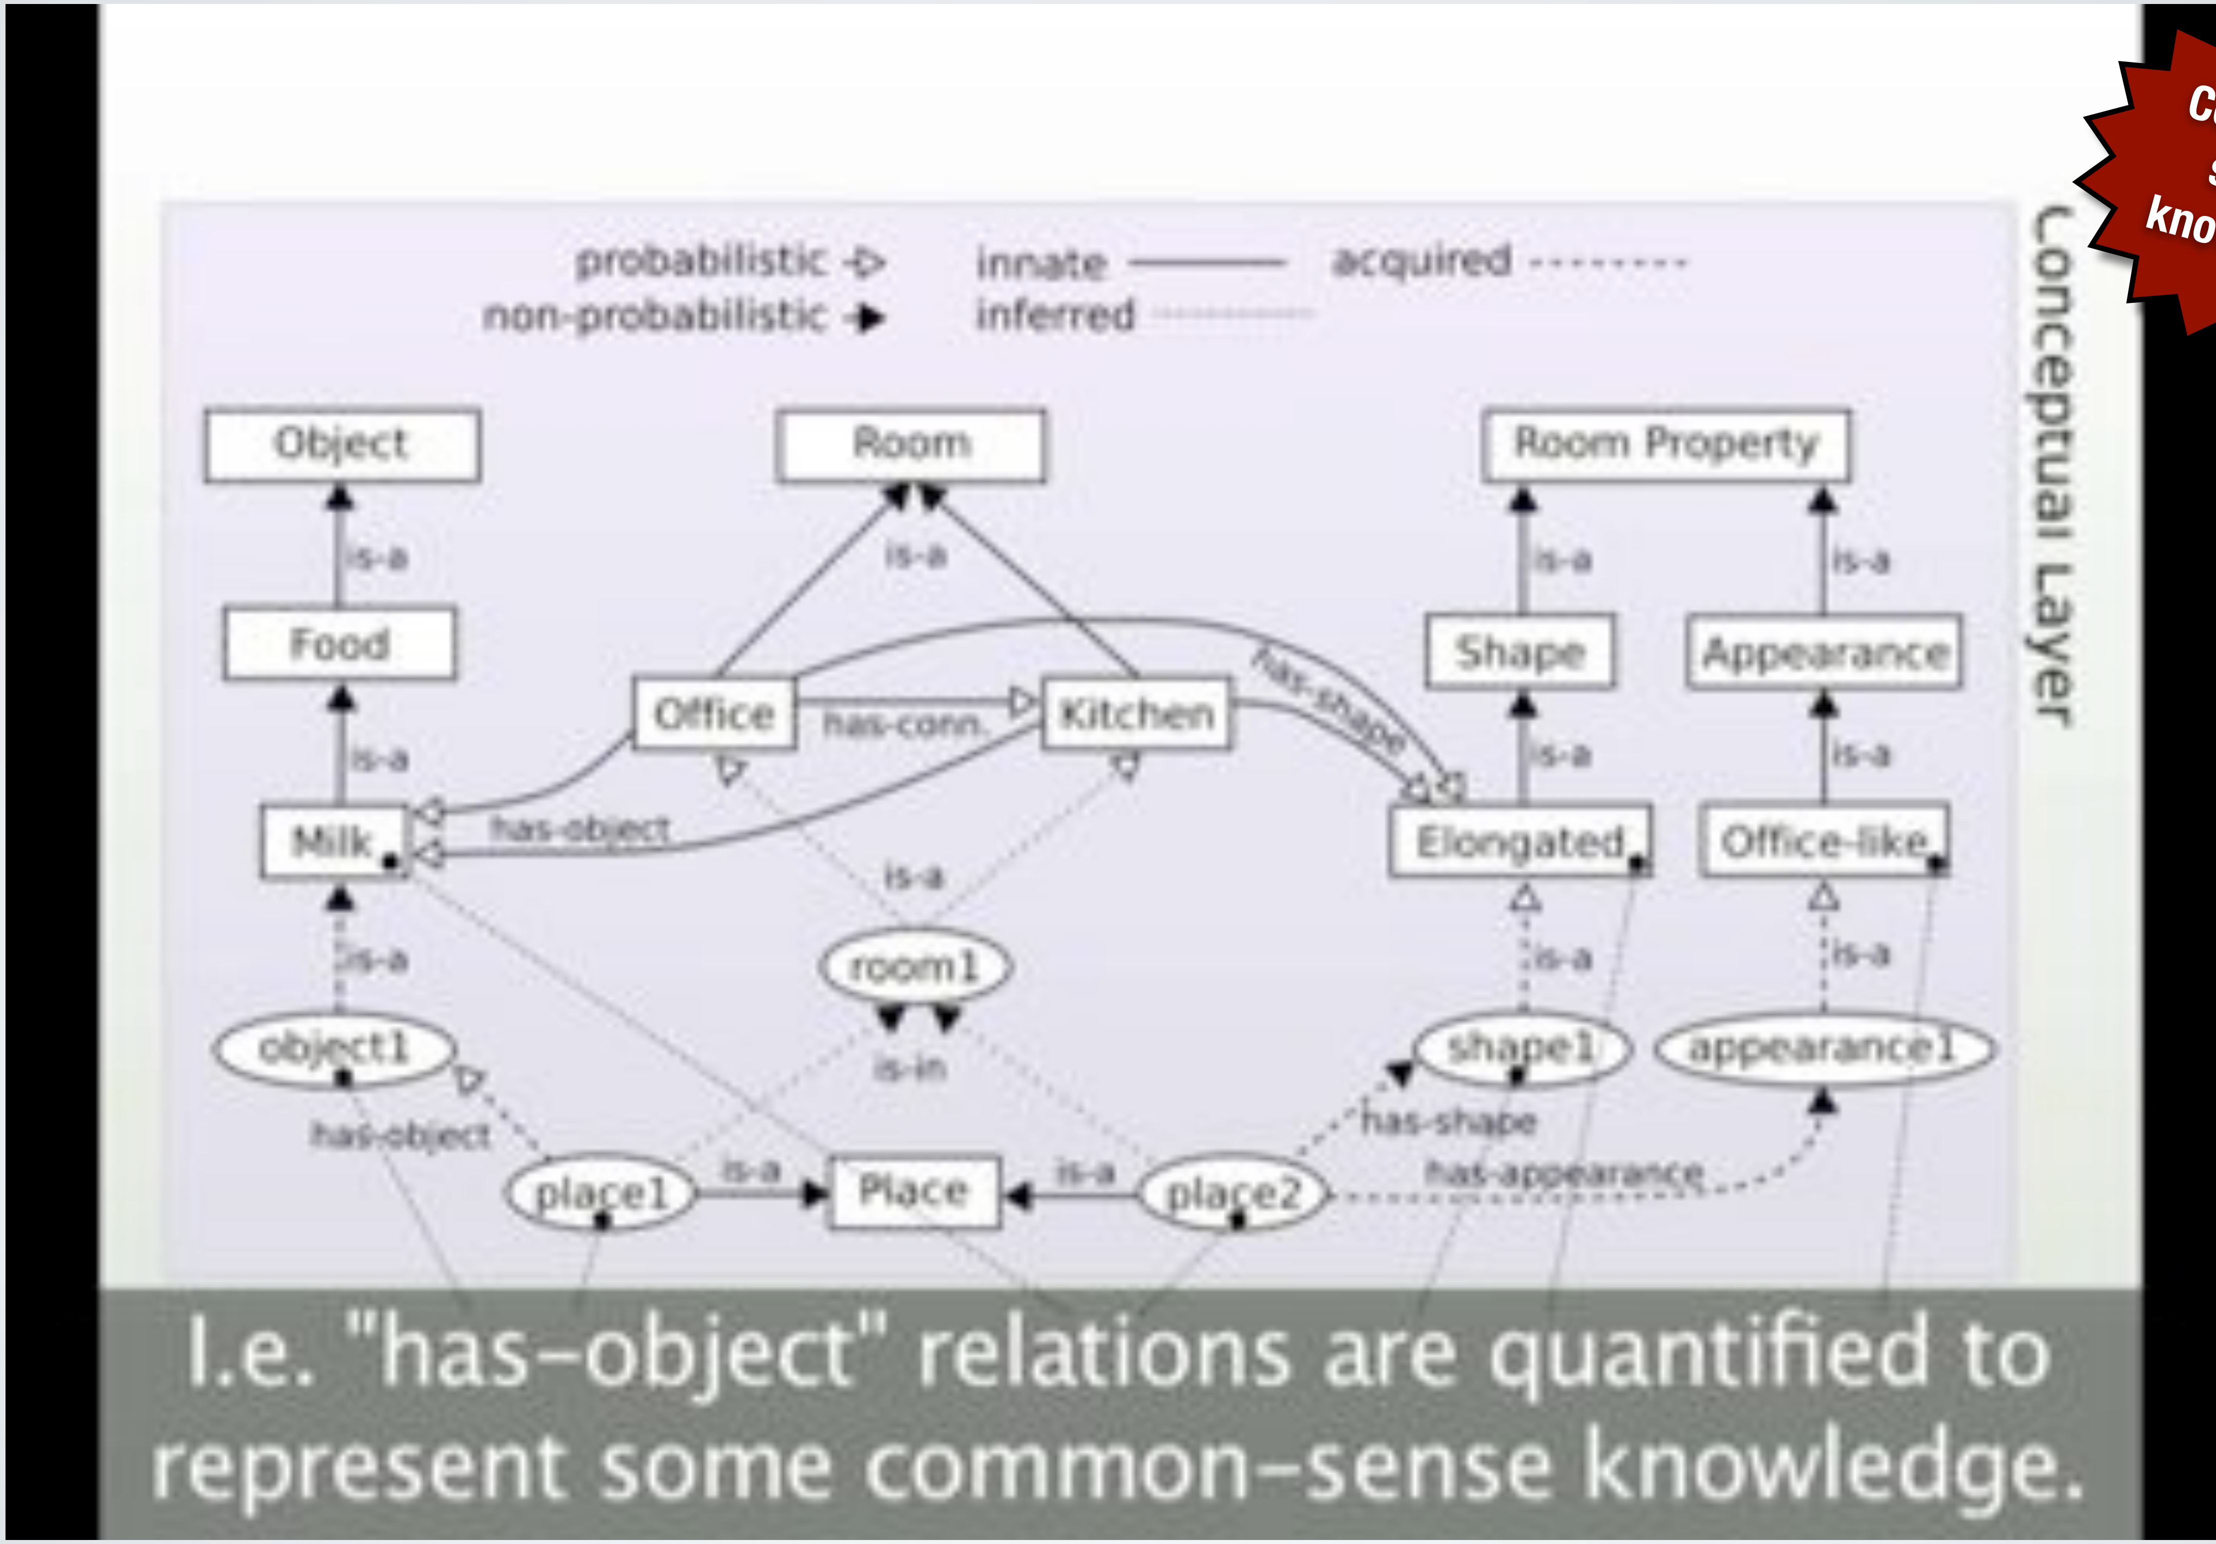

Common-sense knowledge

"HARVESTING" COMMON SENSE

- leverage common-sense knowledge implicitly available in the world wide web

- object and location concepts are "bootstrapped" from the 'locations' database provided by the Open Mind Indoor Common Sense (OMICS) database (Honda Research USA)

- co-occurrence frequency yielding probabilities estimated by counting the number of hits in an image search engine

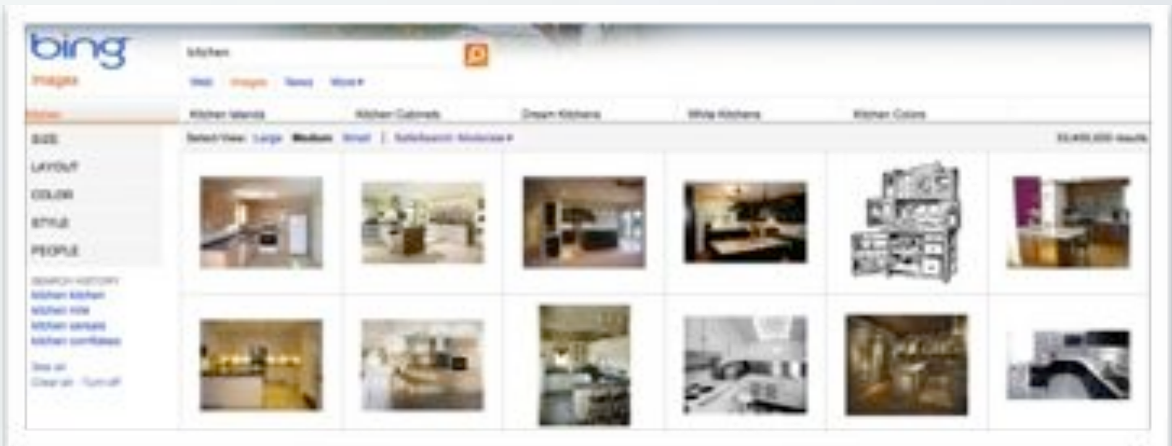

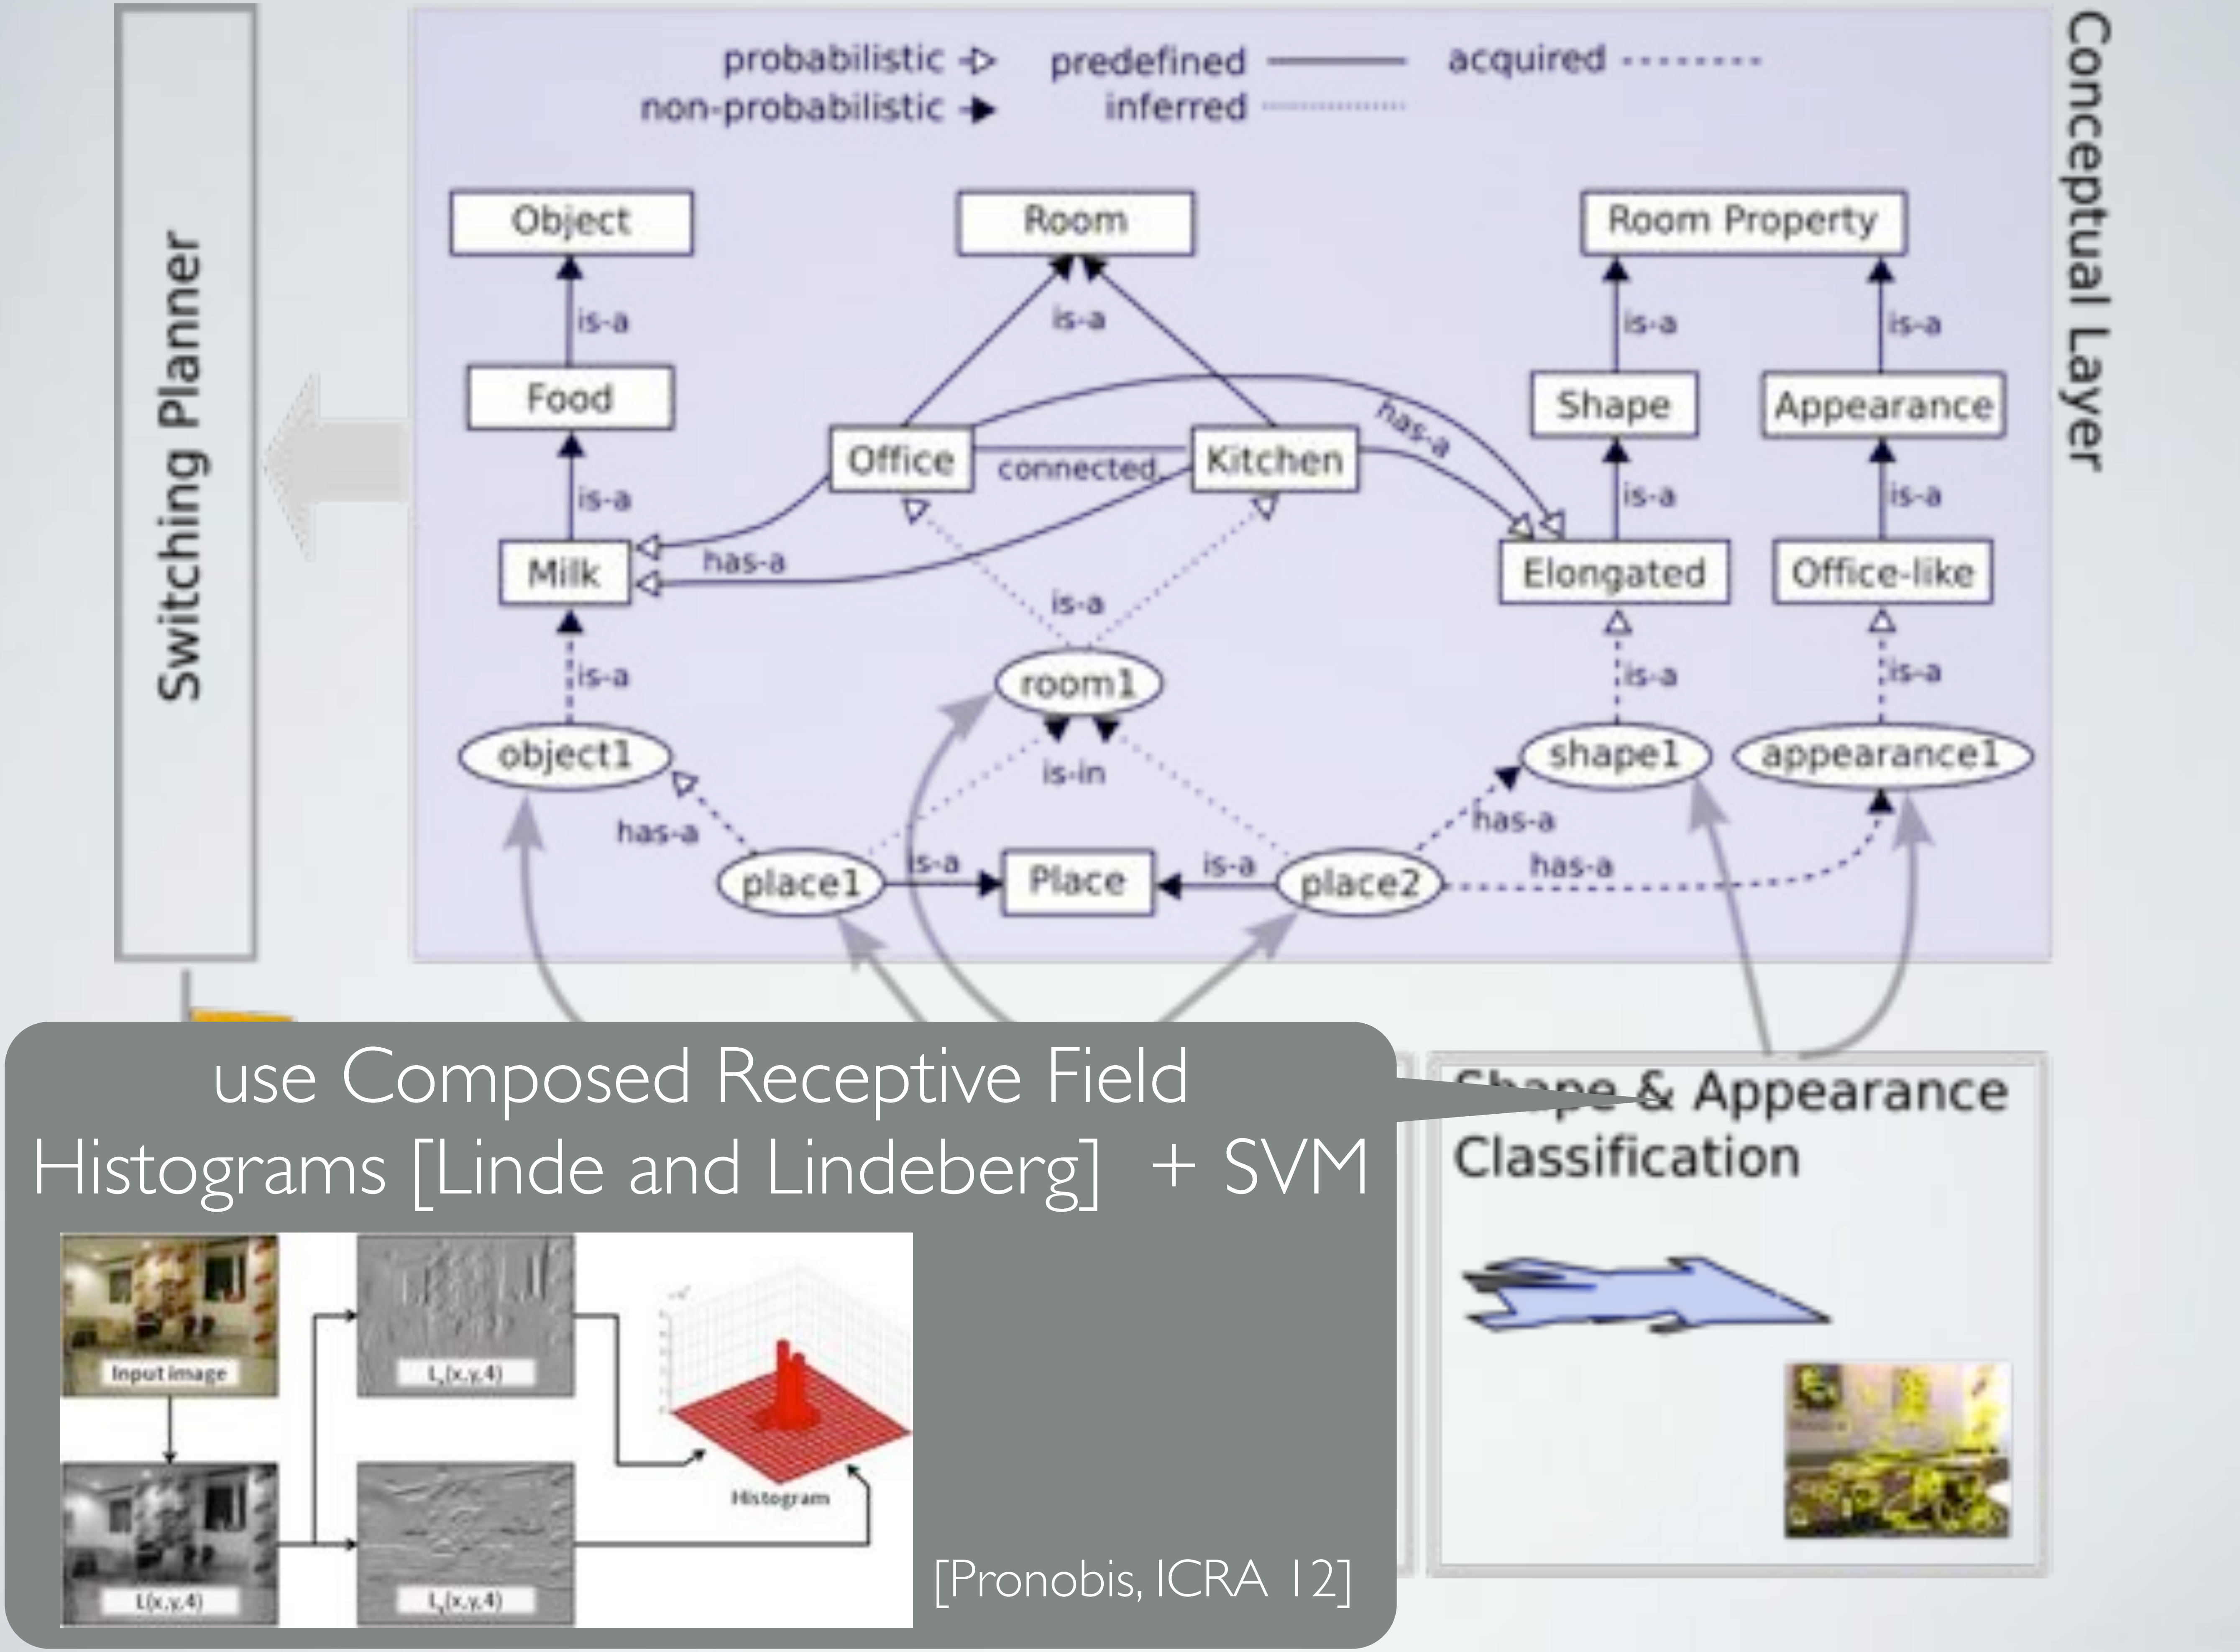
Switching Planner
Conceptual Layer
probabilistic
non-probabilistic
predefined
inferred
acquired
Object
Room
Room Property
Food
Office
Kitchen
Shape
Appearance
Milk
Elongated
Office-like
room1
object1
shape1
appearance1
place1
Place
place2
is-a
has-a
connected
is-in
Shape & Appearance Classification
use Composed Receptive Field Histograms [Linde and Lindeberg] + SVM
Input image
Histogram
[Pronobis, ICRA 12]

CHAIN GRAPH REASONING

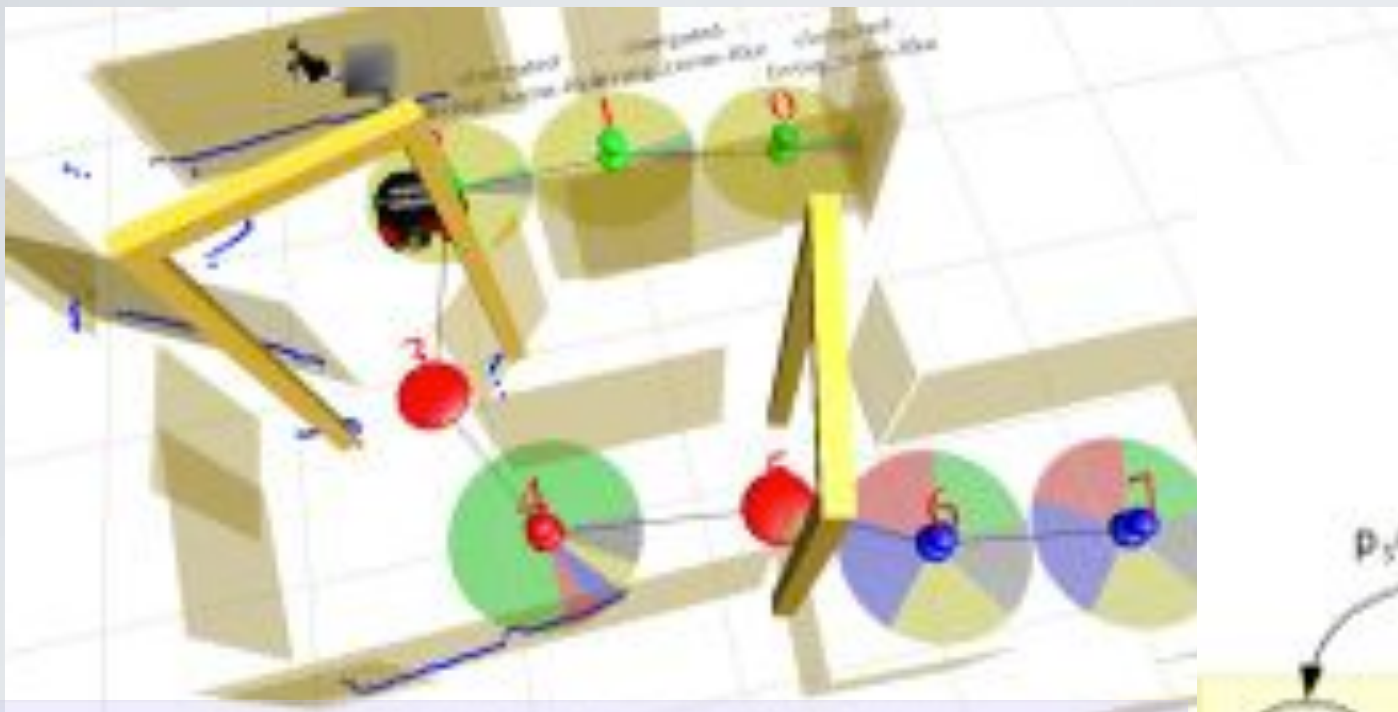

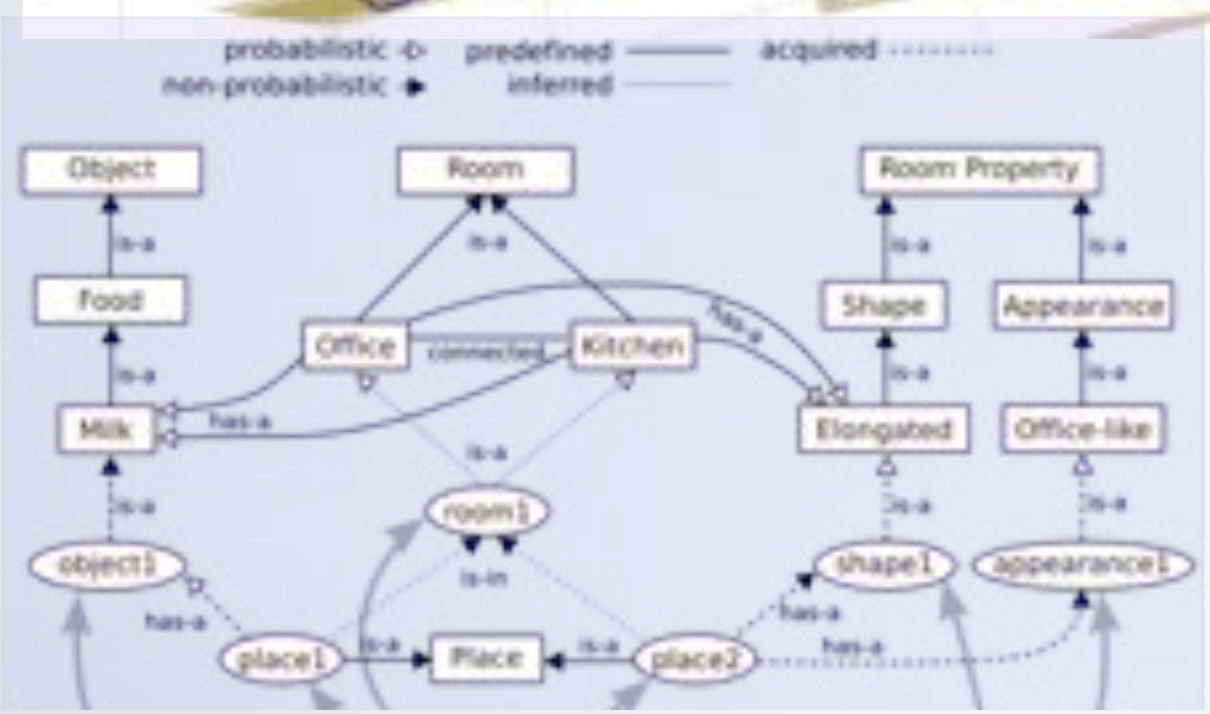

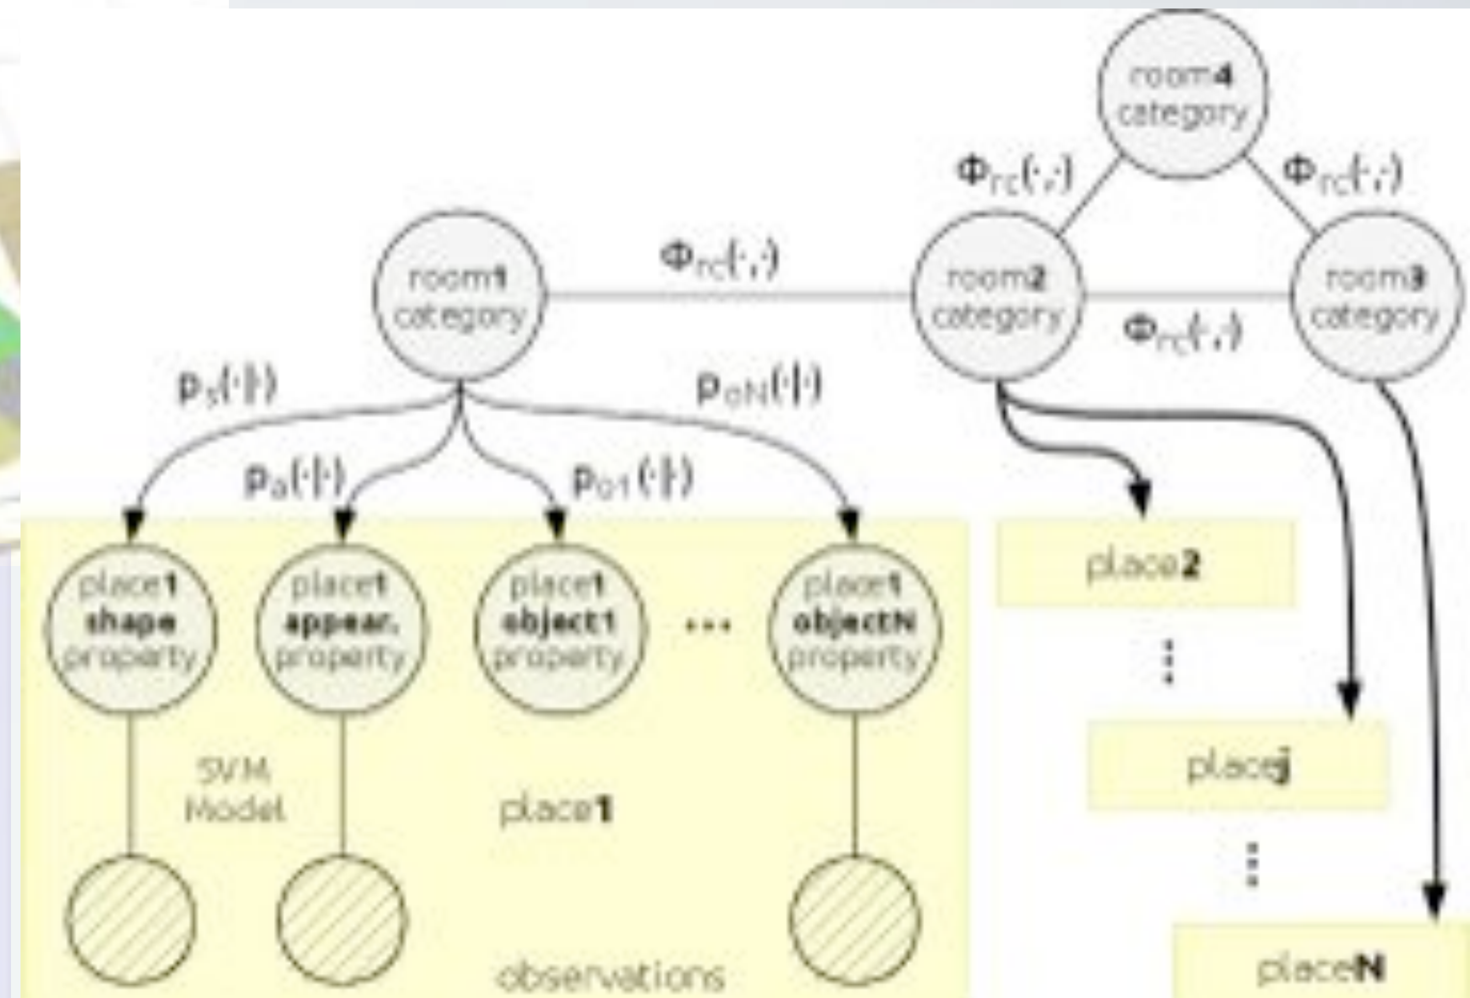

probabilistic belief state

Lauritzen et al: Chain graph models and their causal interpretations,
J. Roy. Statistical Society, 2002

INFER THE ROOM CATEGORY

Room Categorisation

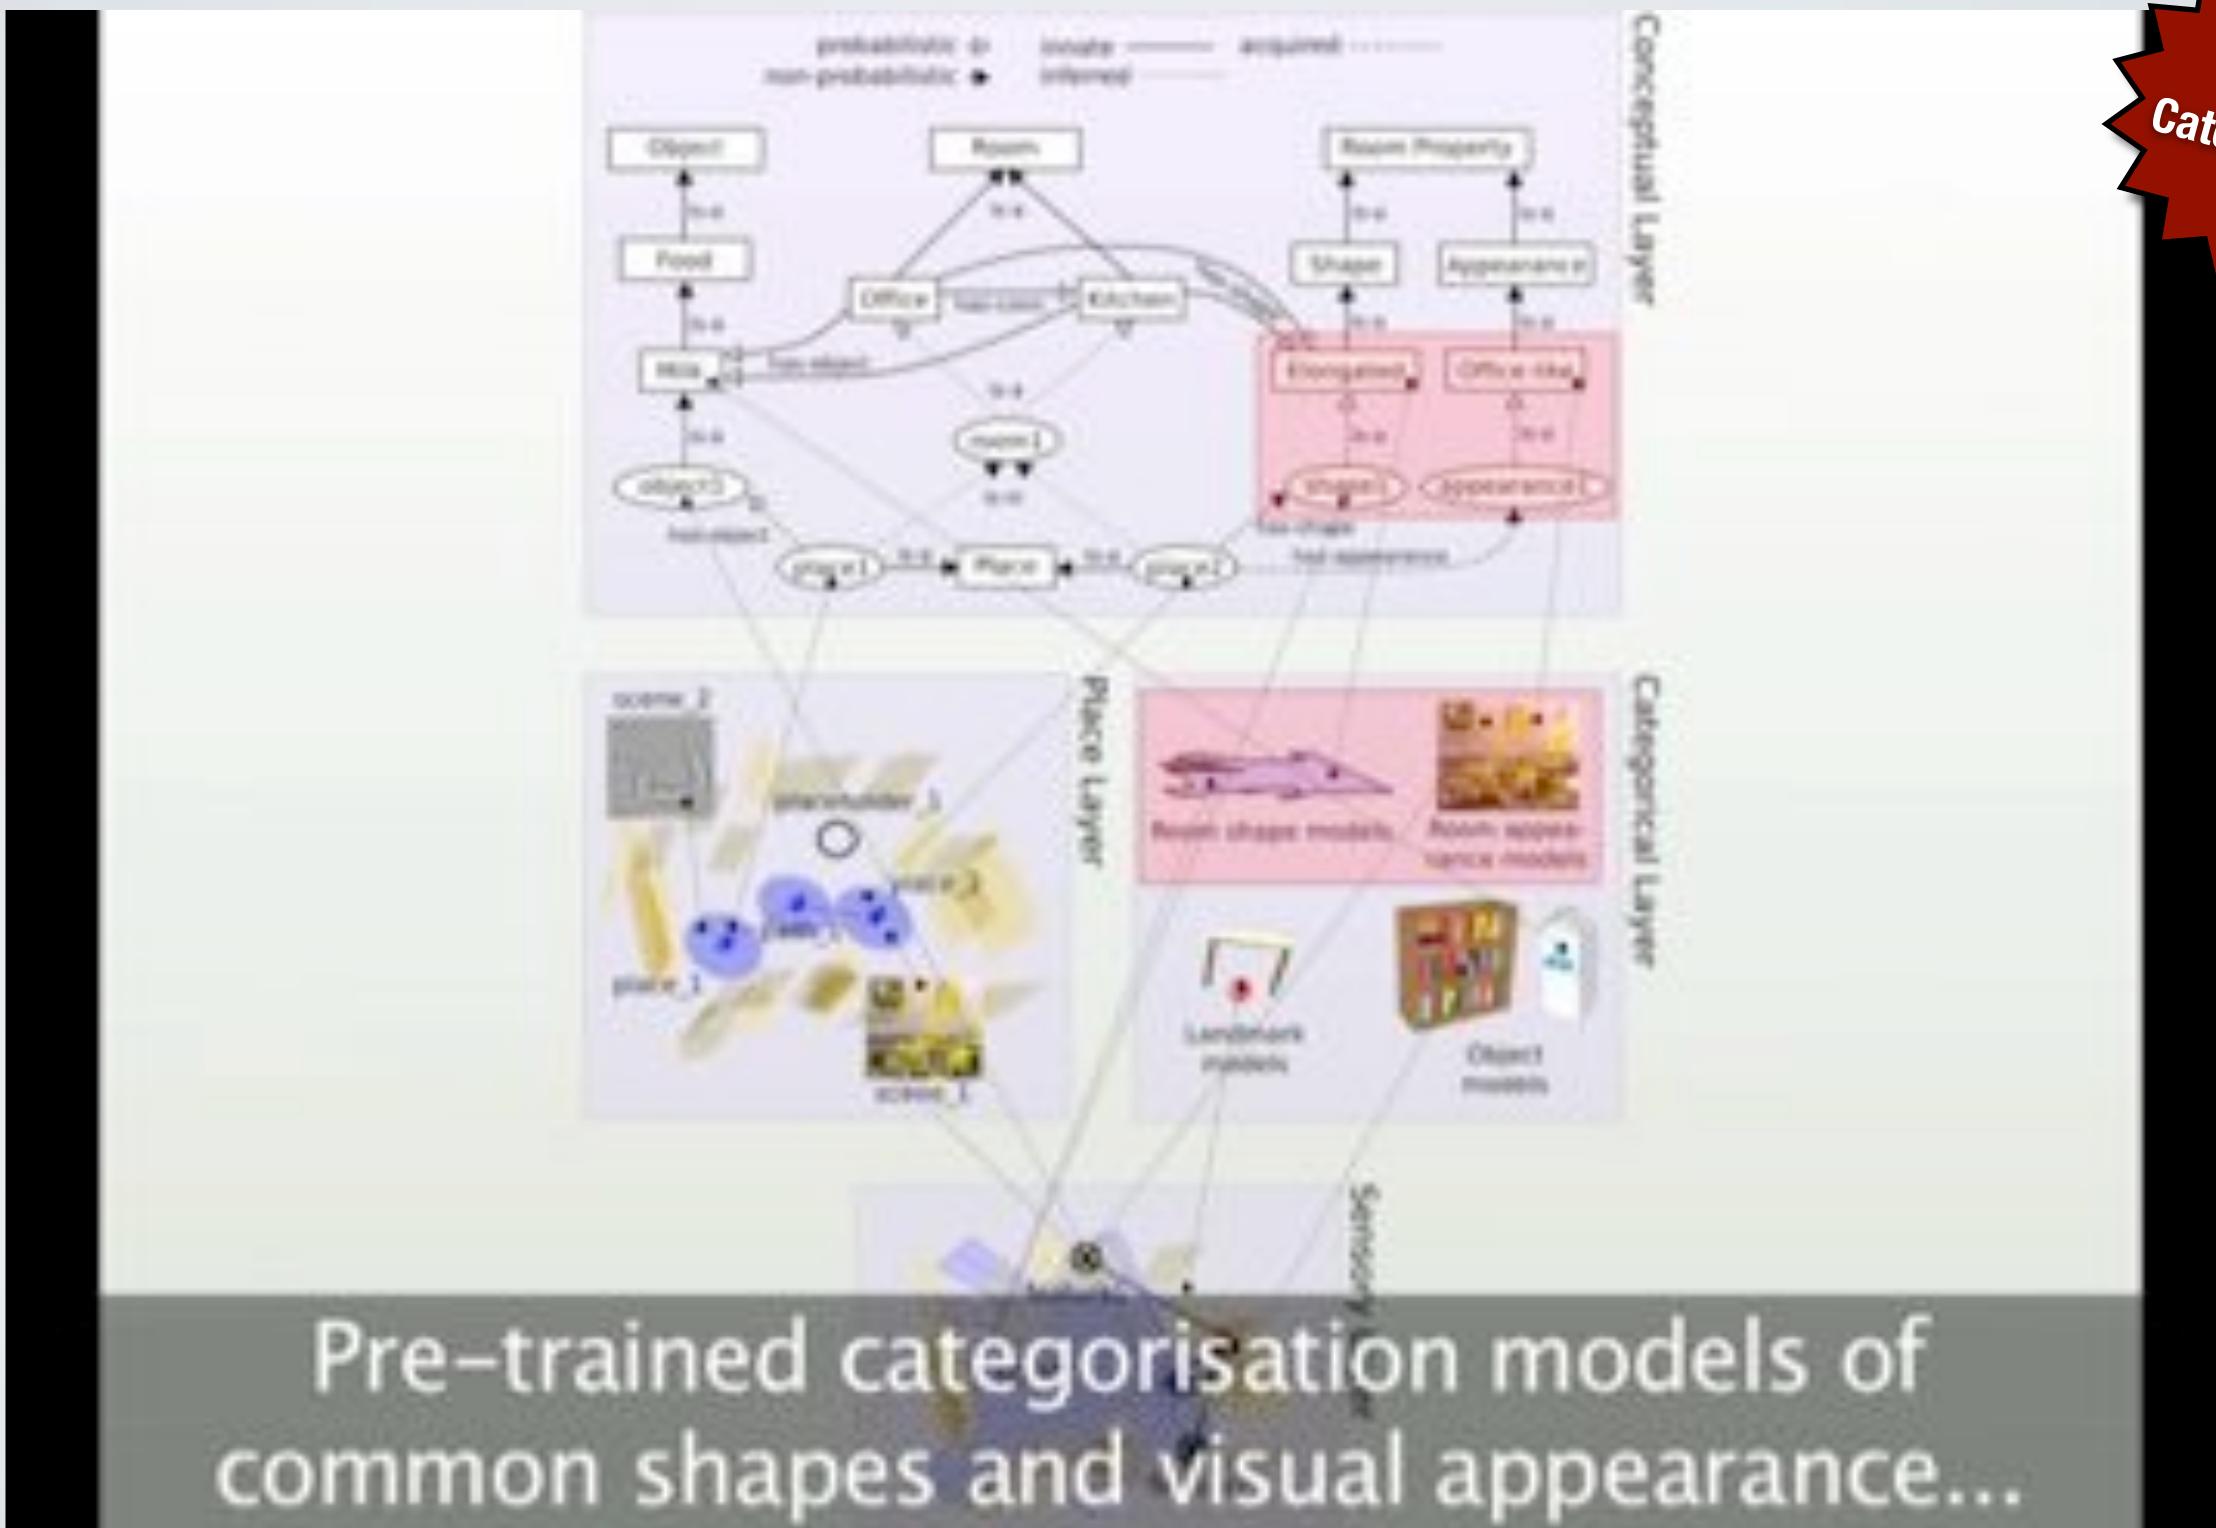

THE COGX APPROACH

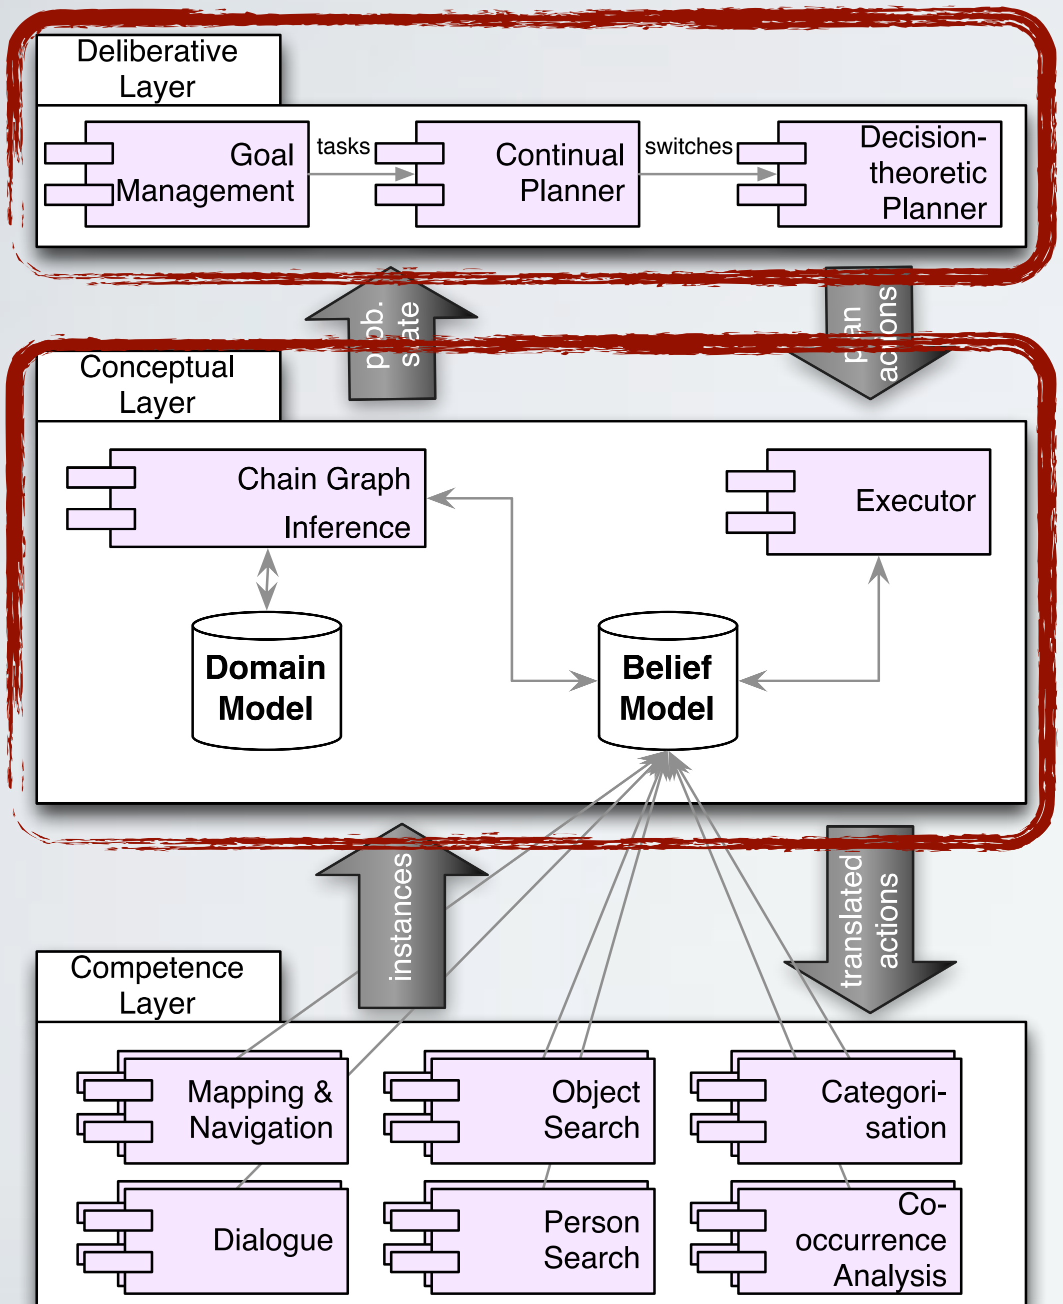

- domain-independent deliberative layer
  - plan autonomous behaviour to accomplish various goals in complex, probabilistic state spaces
- domain-dependent conceptual layer
  - accommodates probabilistic domain and instance knowledge
- competence layer
  - implements sensing and acting
- implemented using CAST (not YARP)

http://www.cs.bham.ac.uk/research/projects/cosy/cast/

PLANNING UNDER UNCERTAINTY

Probabilistic Planning

- world state cannot be observed directly (uncertain sensing)
- "naturally" represented as POMDP
- probabilistic state space can be intractably huge (here $\sim 10^{27}$)
- domain-independent solution: *switching planner*
  - make *assumptions* in a *sequential* session to generate most rewarding trace
  - *contingent* session for observation planning
  - (dis)confirmation of assumptions by contingent session

Hanheide et al., IJCAI 2011
Göbelbecker et al., AAAI 2011

SWITCHING PLANER

- *Continual* (classical) planner: good performance, given useful *assumptions*
- *Decision-theoretic* POMDP planner: very good at reasoning about how observations provide evidence about the underlying world state
- *Switch* between sessions to exploit both's strengths

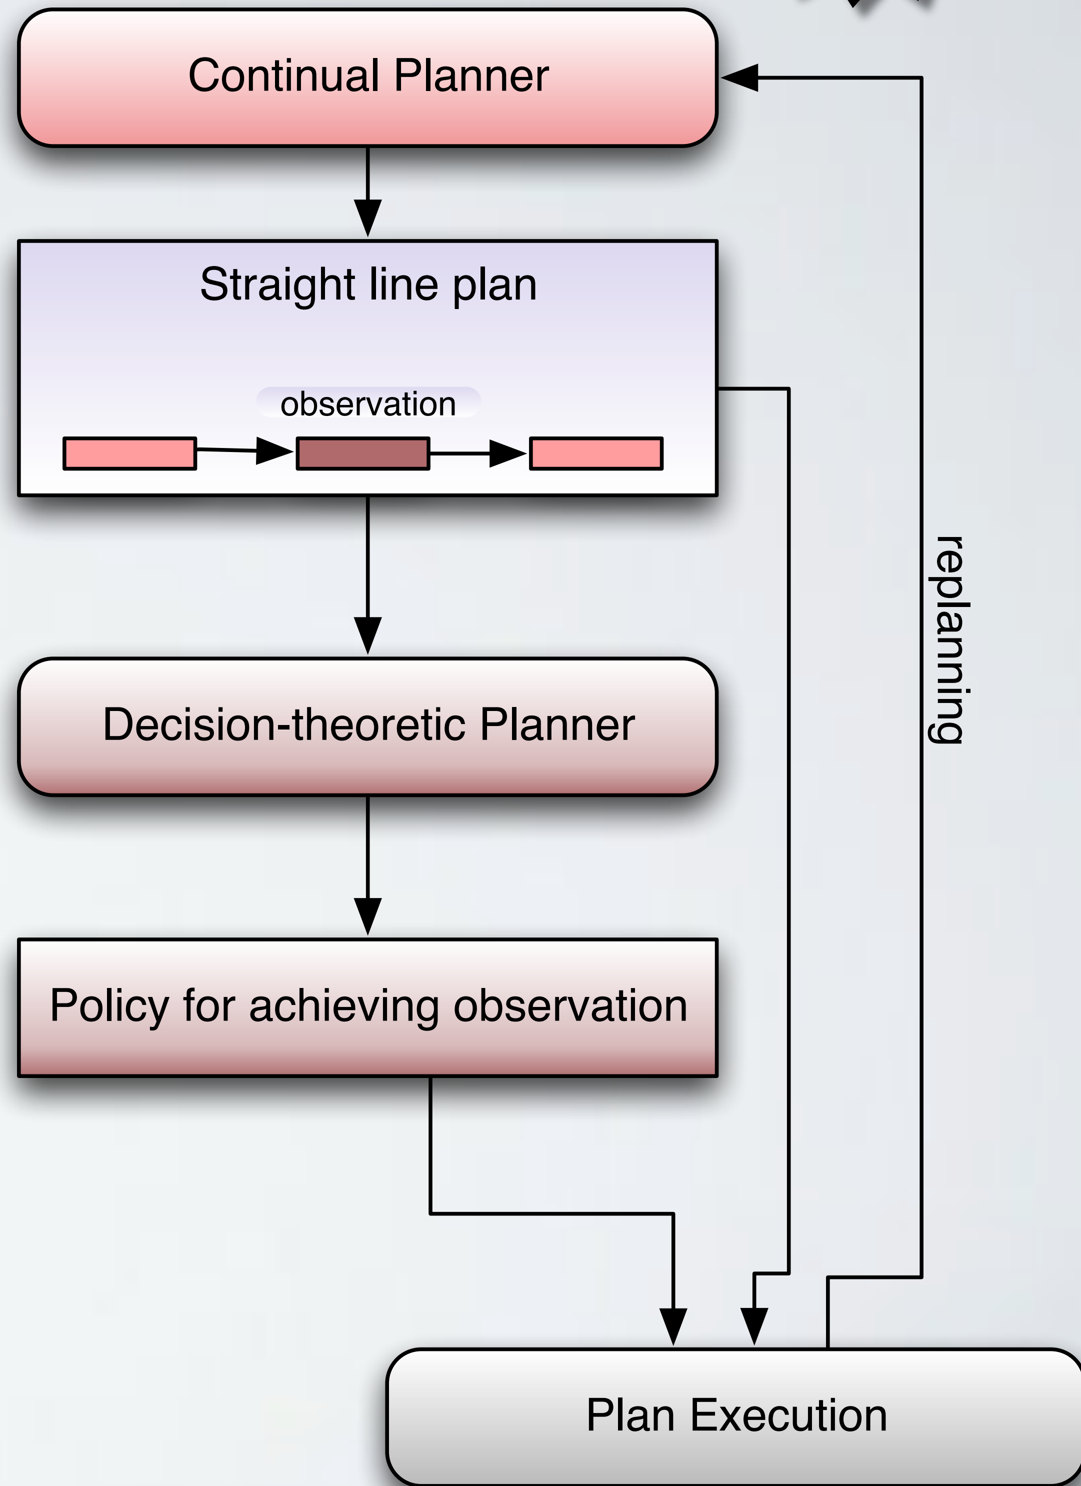

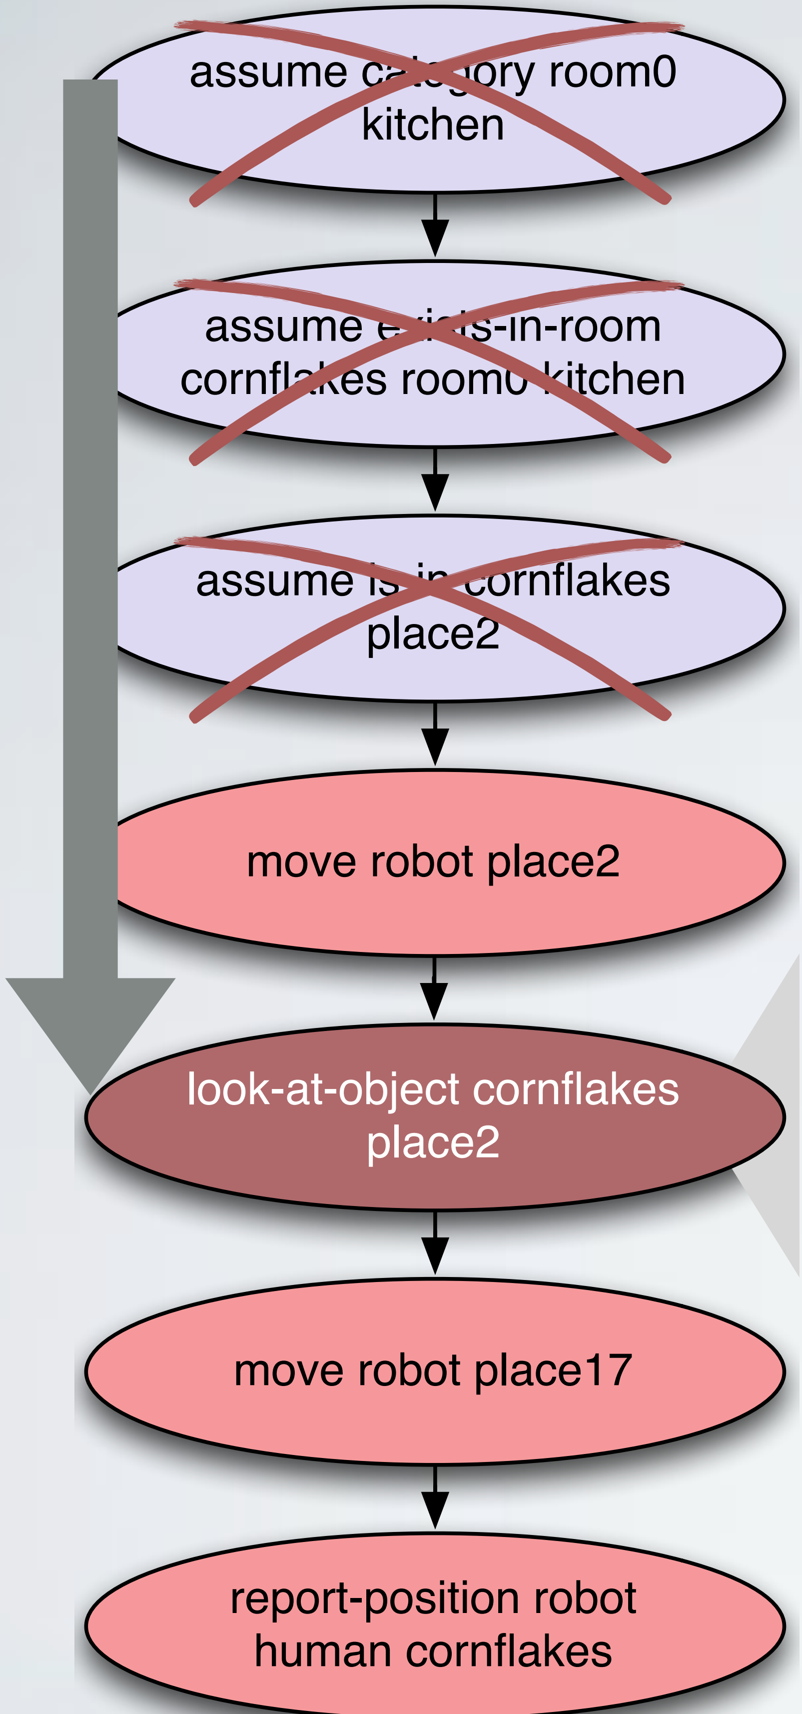

- execute sequential plan until there is a sensing action

- invoke DT planner on abstract problem by retracting assumptions

- assumptions are retracted in ascending order of the conditional entropy

$$\mathrm{H}(X|y) = \sum_{x \in \mathcal{X}, y' \in \{y, \neg y\}} p(x \wedge y') \log_2 \frac{p(y')}{p(x \wedge y')}$$

with y being the assumption, and X the set of propositions in the abstract DT process

- retraction of deterministic assumptions stops when N(=150) states are true with non-zero probability

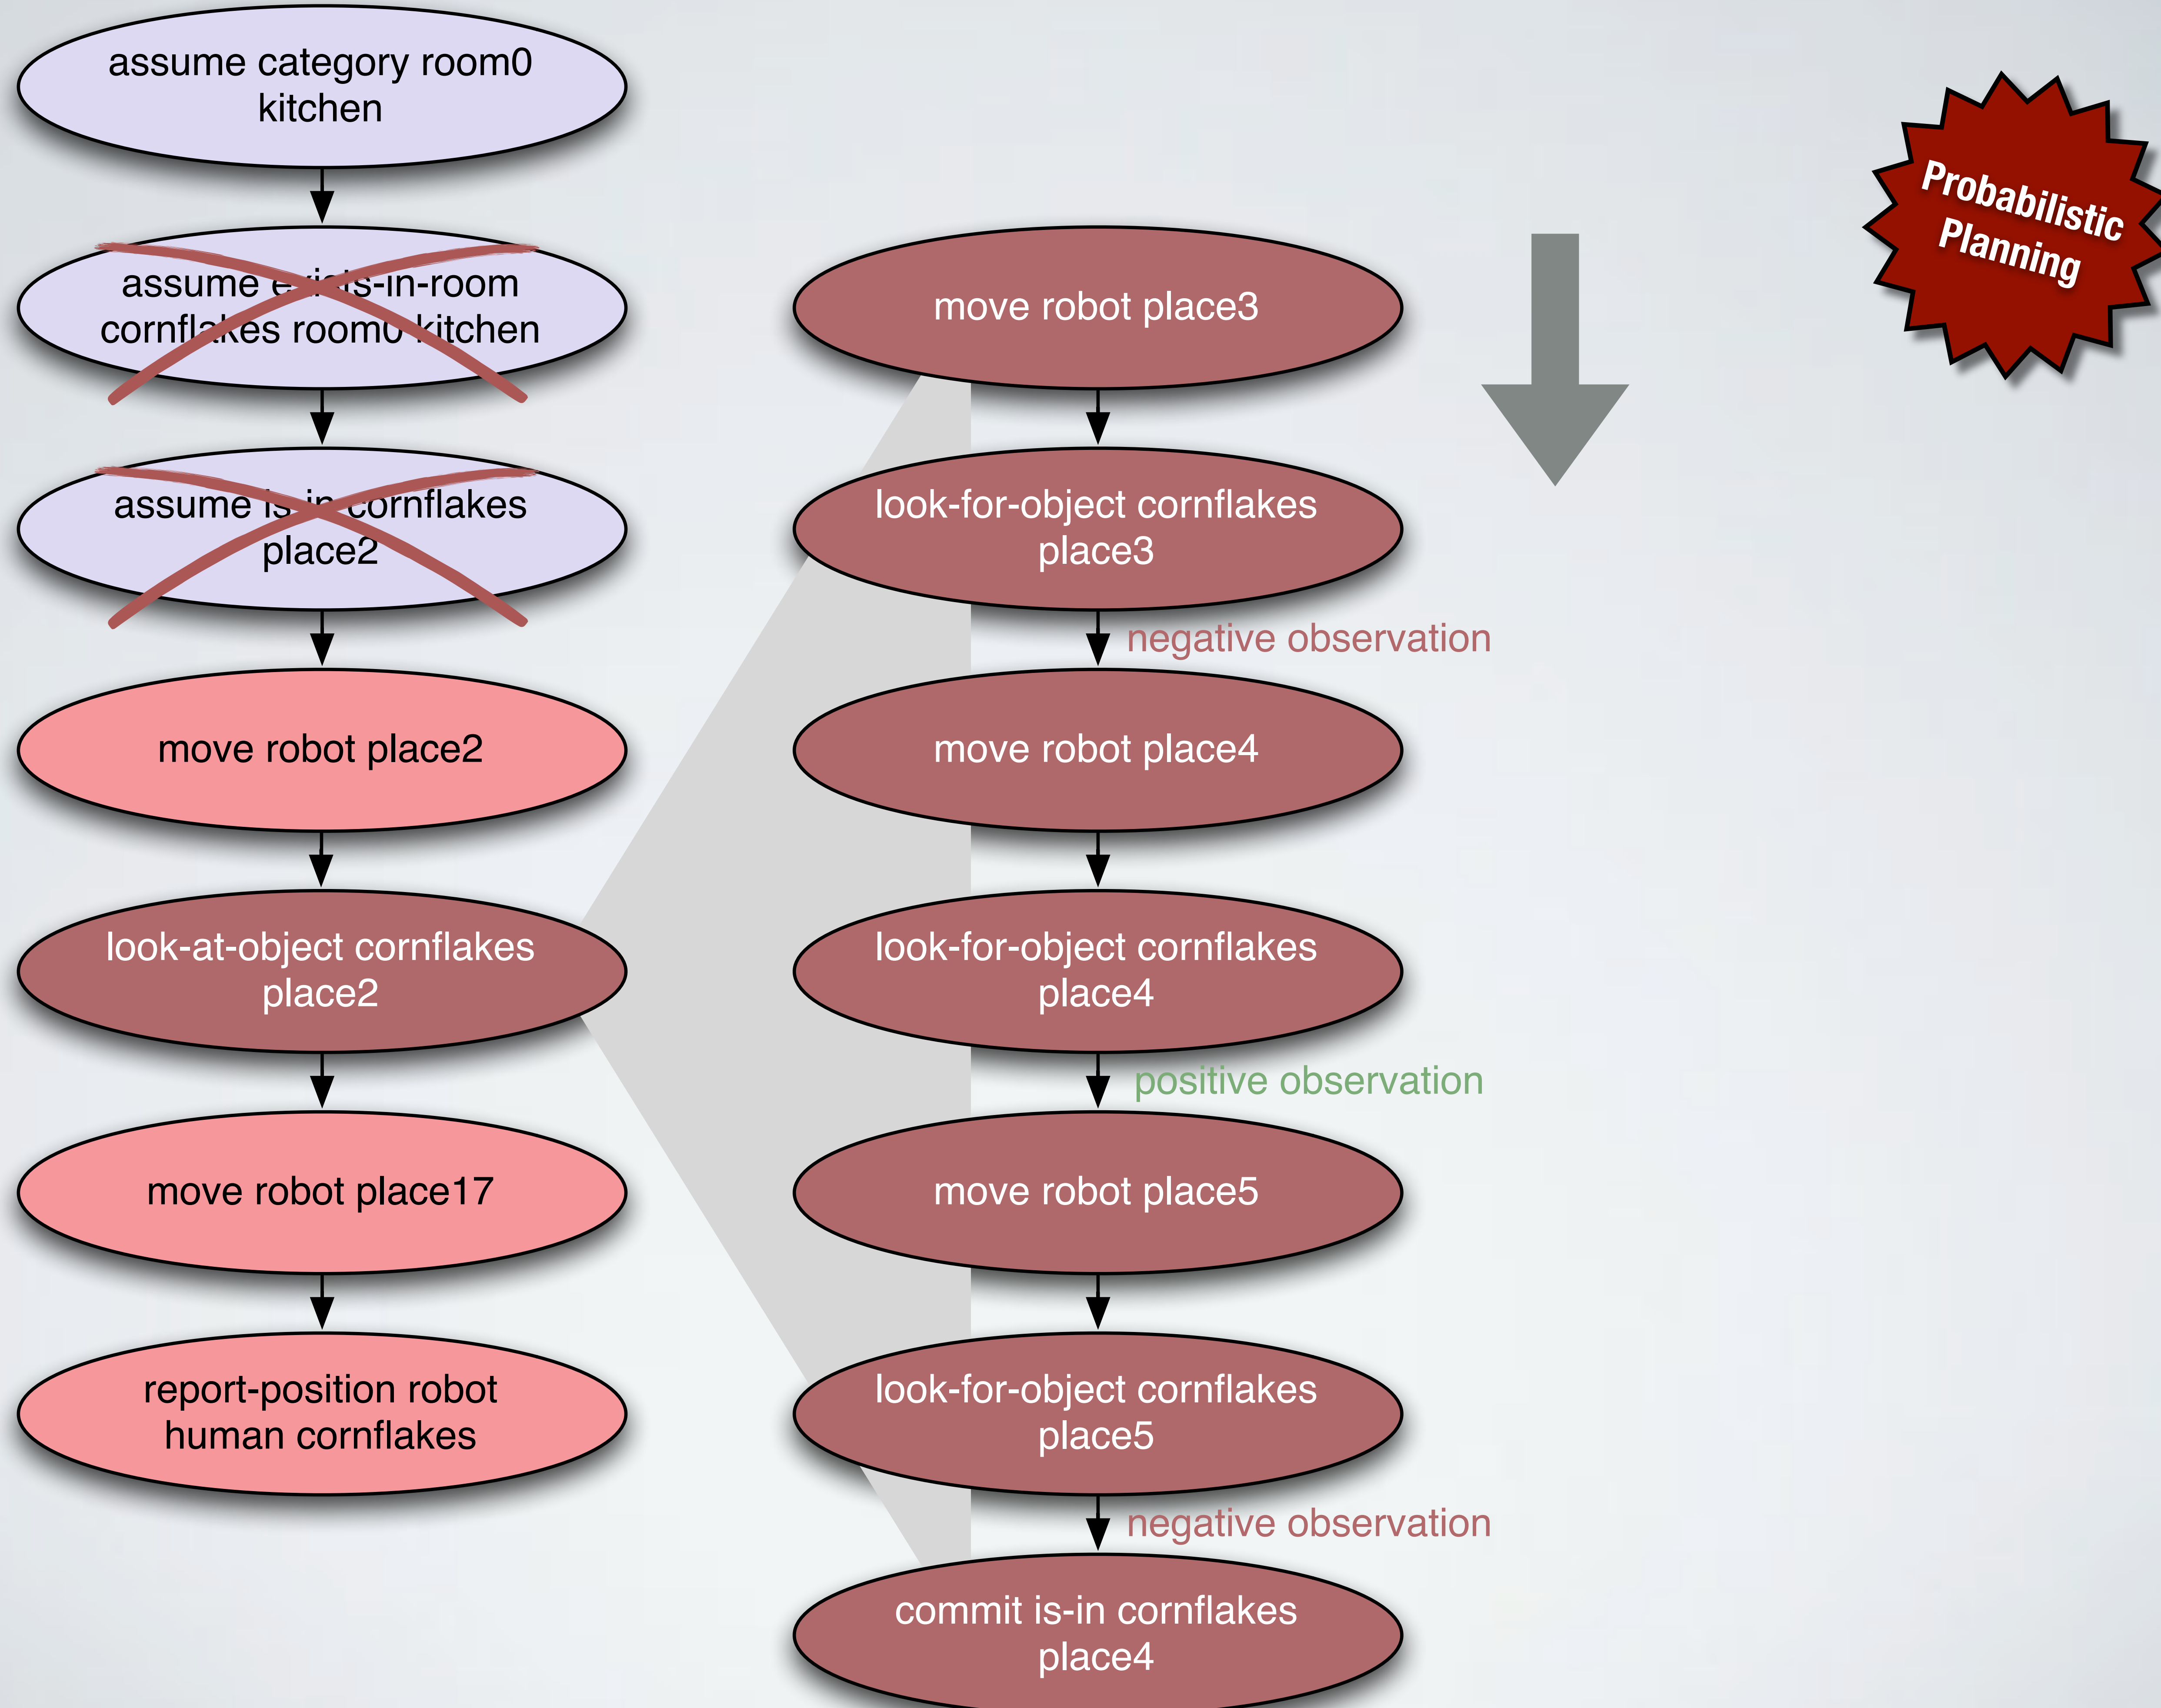

assume category room0 kitchen
~~assume exists-in-room cornflakes room0 kitchen~~
~~assume is-in cornflakes place2~~
move robot place2
look-at-object cornflakes place2
move robot place17
report-position robot human cornflakes
move robot place3
look-for-object cornflakes place3
negative observation
move robot place4
look-for-object cornflakes place4
positive observation
move robot place5
look-for-object cornflakes place5
negative observation
commit is-in cornflakes place4
Probabilistic Planning

IF... THEN!

Autonomous Visual Search

... the object is not at (X,Y,Z)? ✓

... the map would be unknown (no global coordinates)?

Autonomous Exploration

... sensing would be uncertain? ✓

Probabilistic Planning

Room Categorisation

... the robot wouldn't know where the shop/kitchen is? ✓

... the common-sense knowledge would be wrong? (✓)

Interactive Hypothesis Verification

Hanheide, M. et al., 2011. Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour. In Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI).

acquire a large map "by hand" prior to any useful services by the robot?

better reason in an "open" world and allow to self-extend

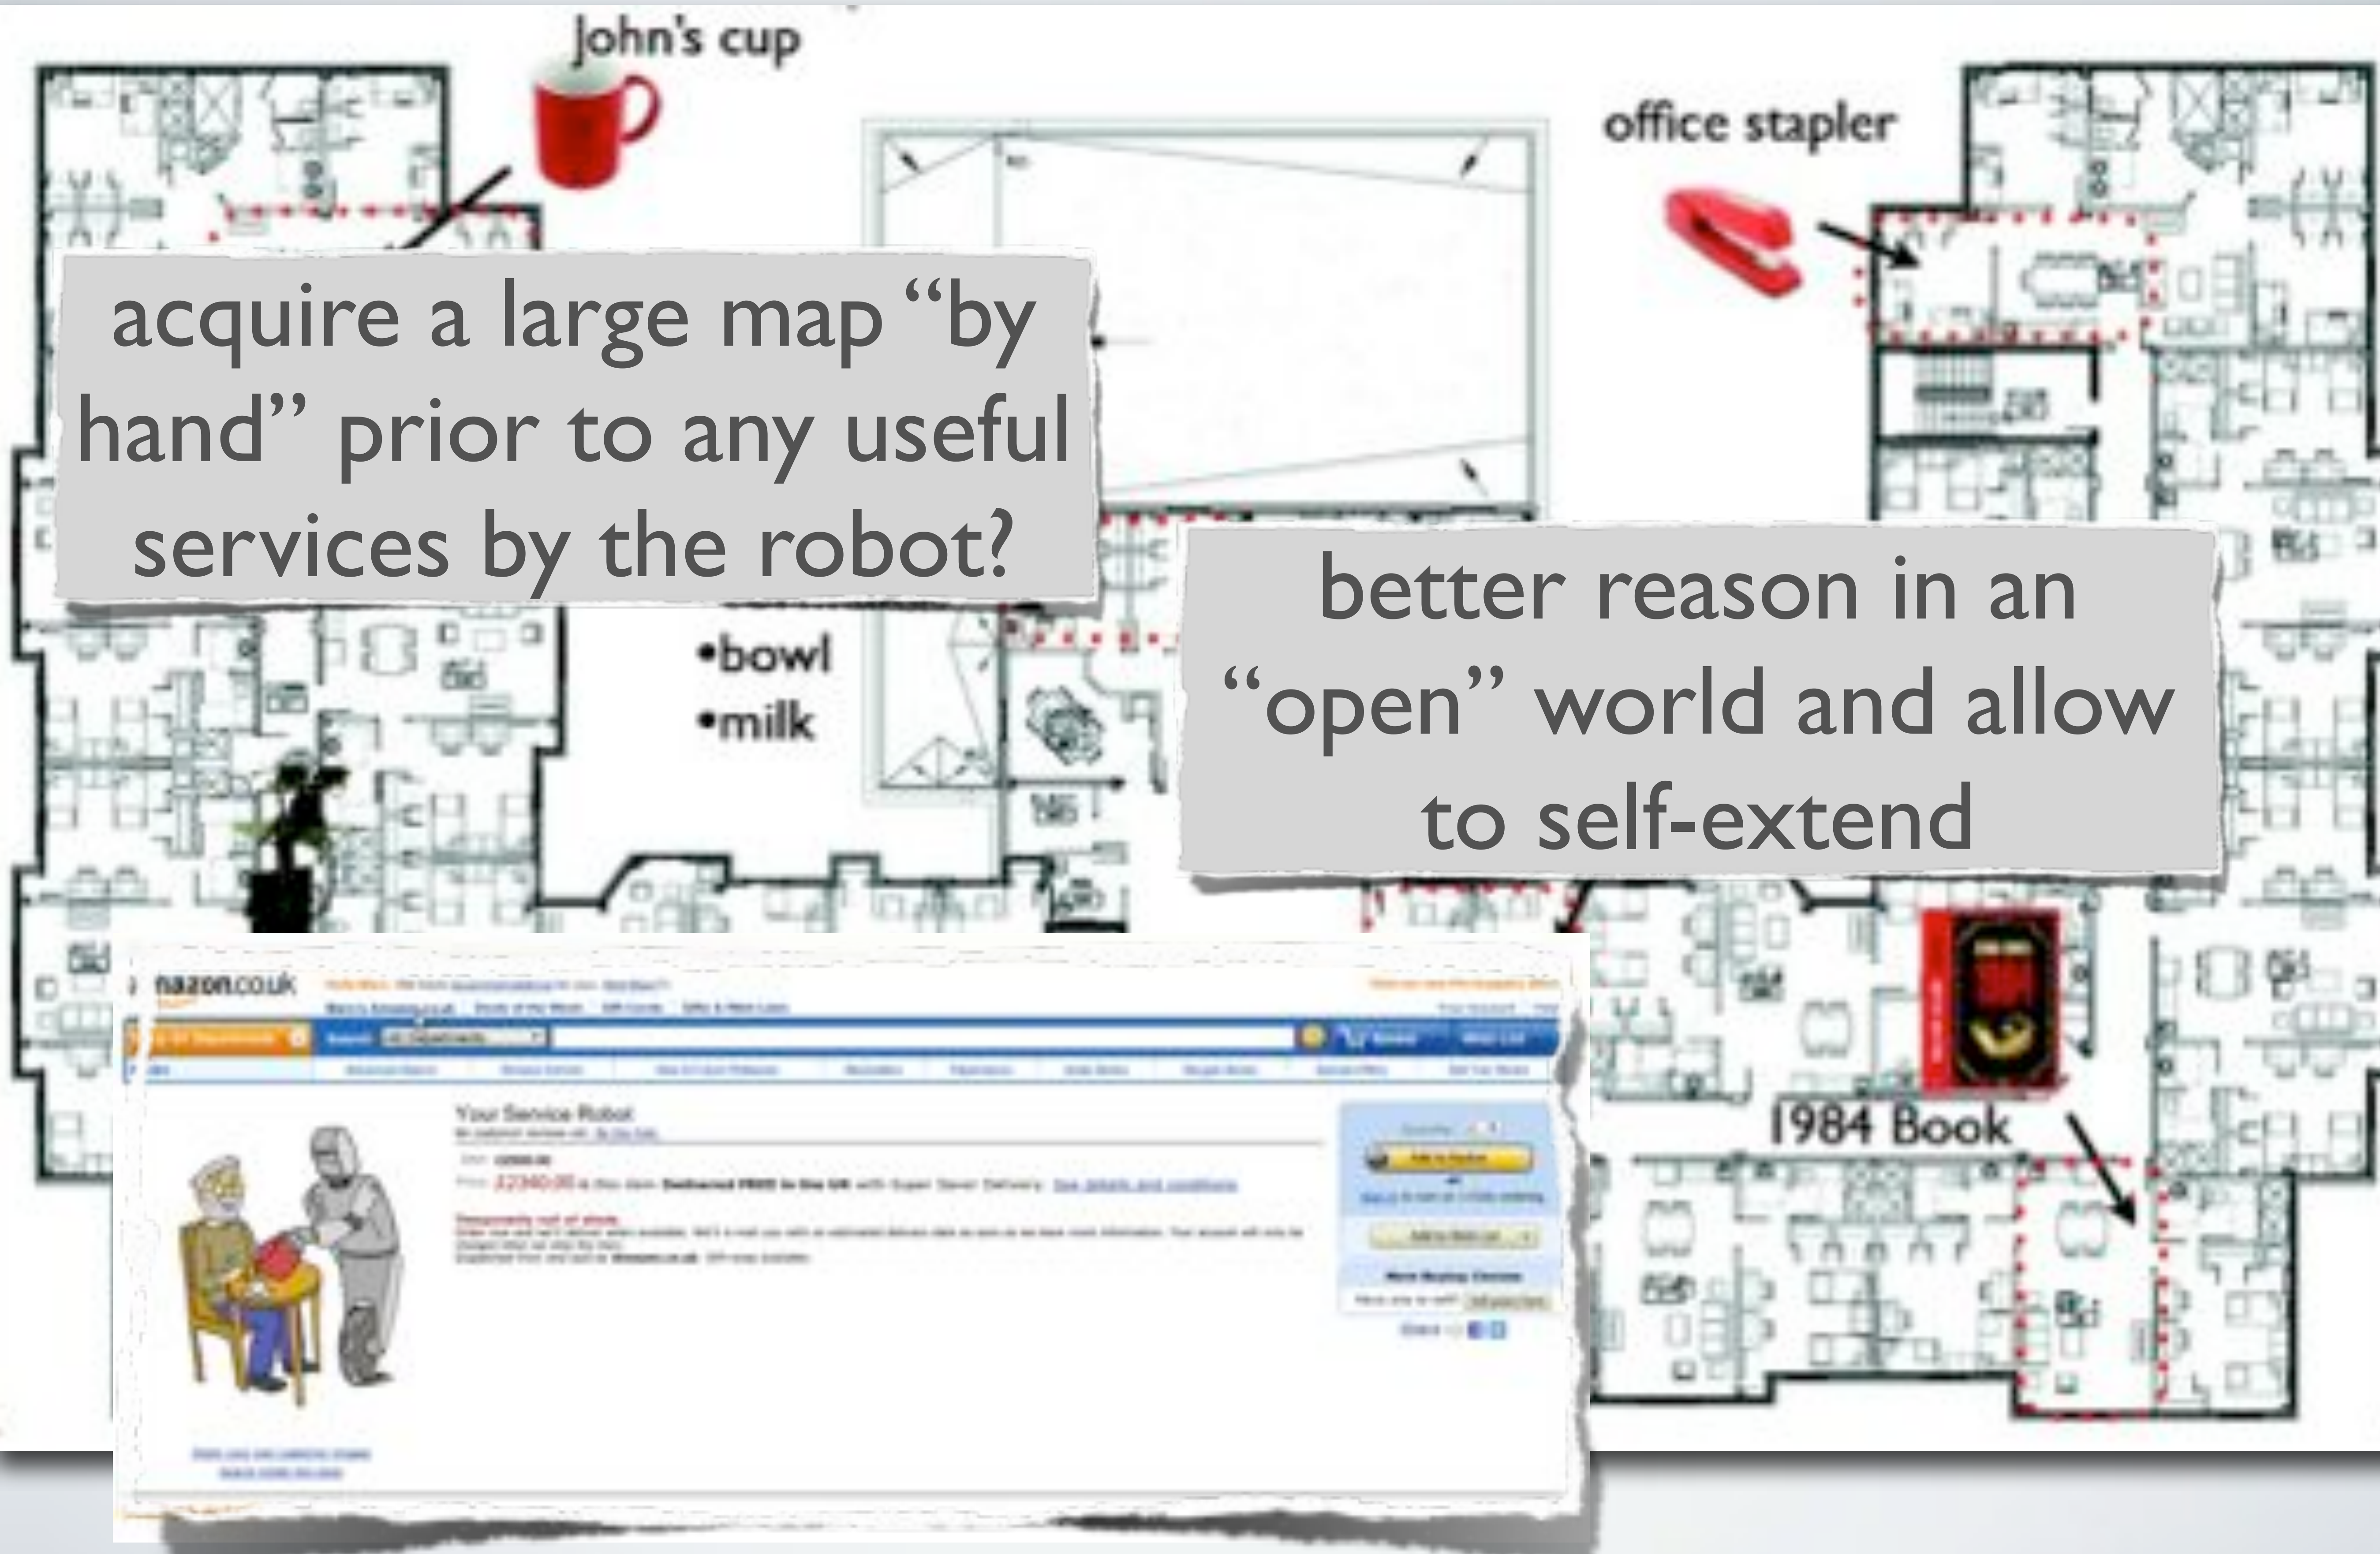

EXPLORATION

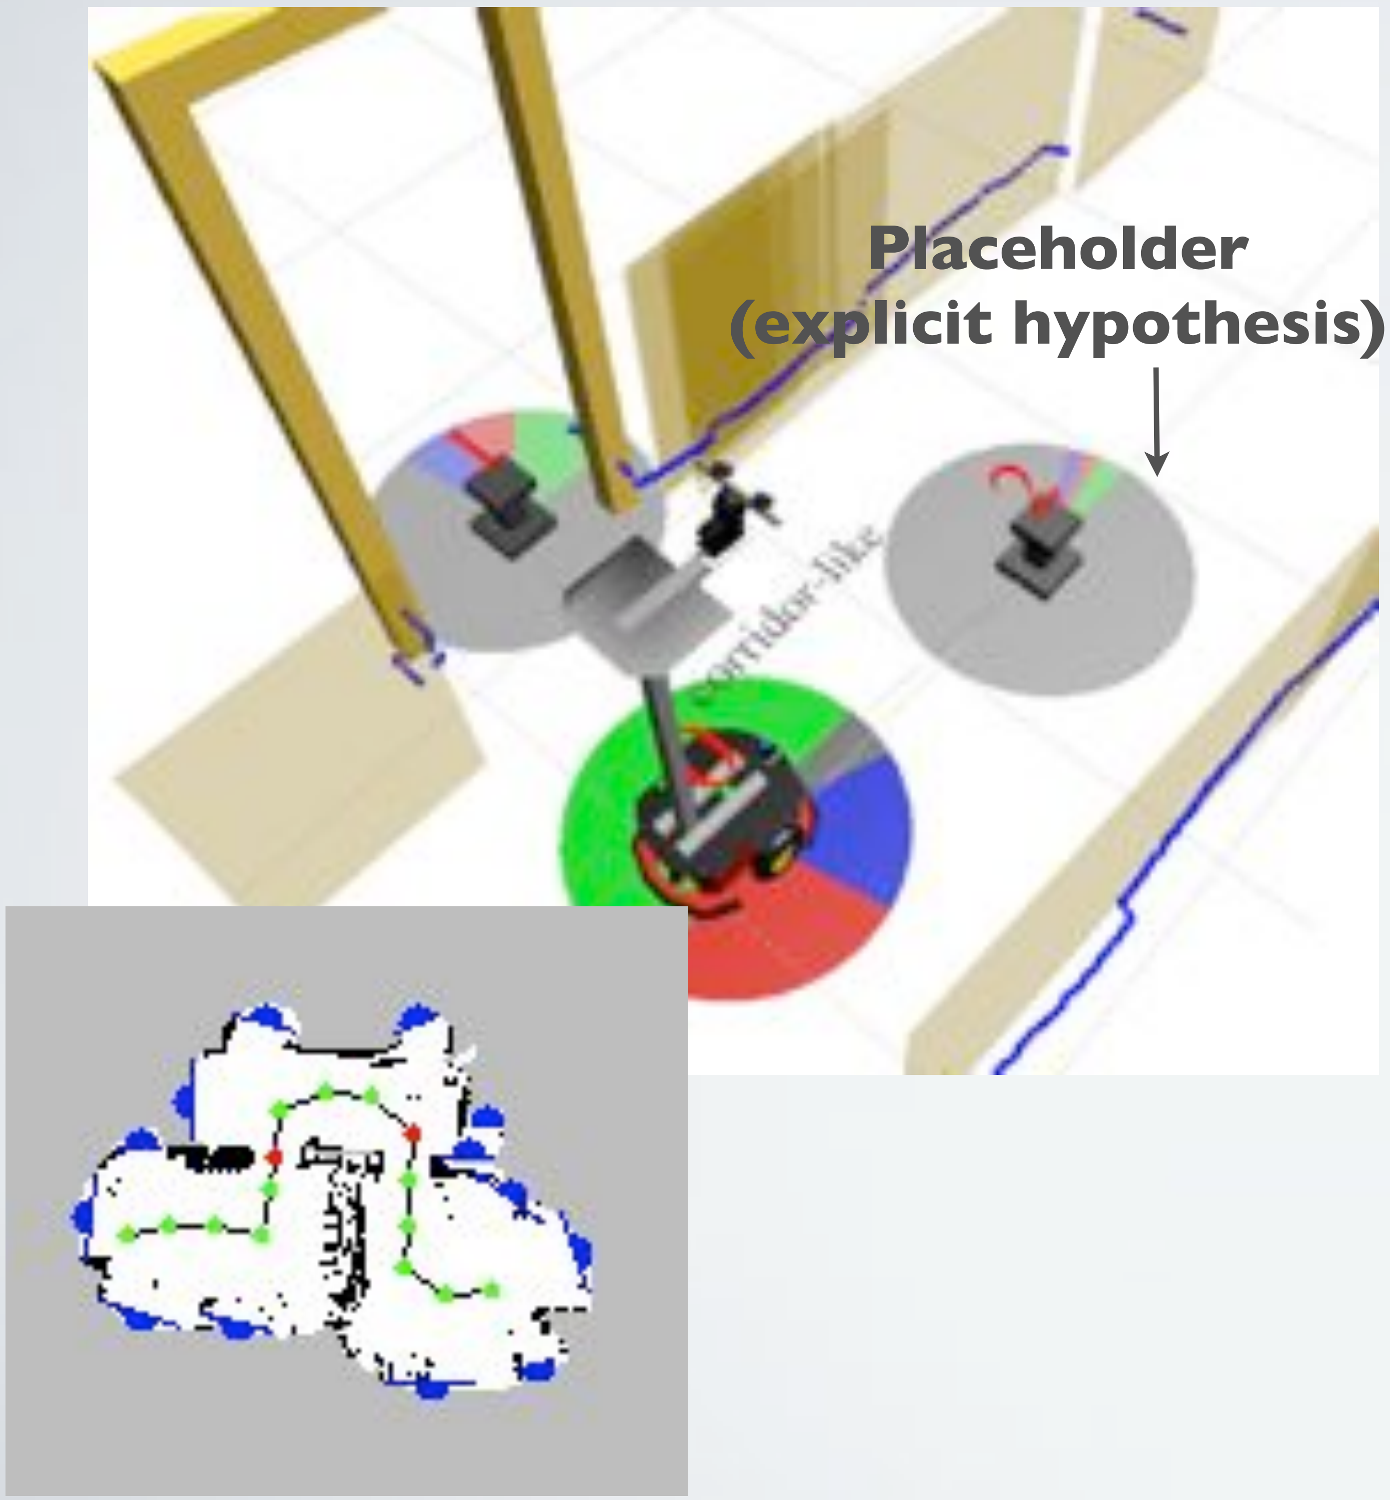

- active visual search in unknown larger scale space
- exploit existing knowledge or explore new hypotheses generated from an analysis of the local metric map (frontiers)
- rooms yield probabilities to find objects
- informed, task-driven exploration or curiosity-driven exploration

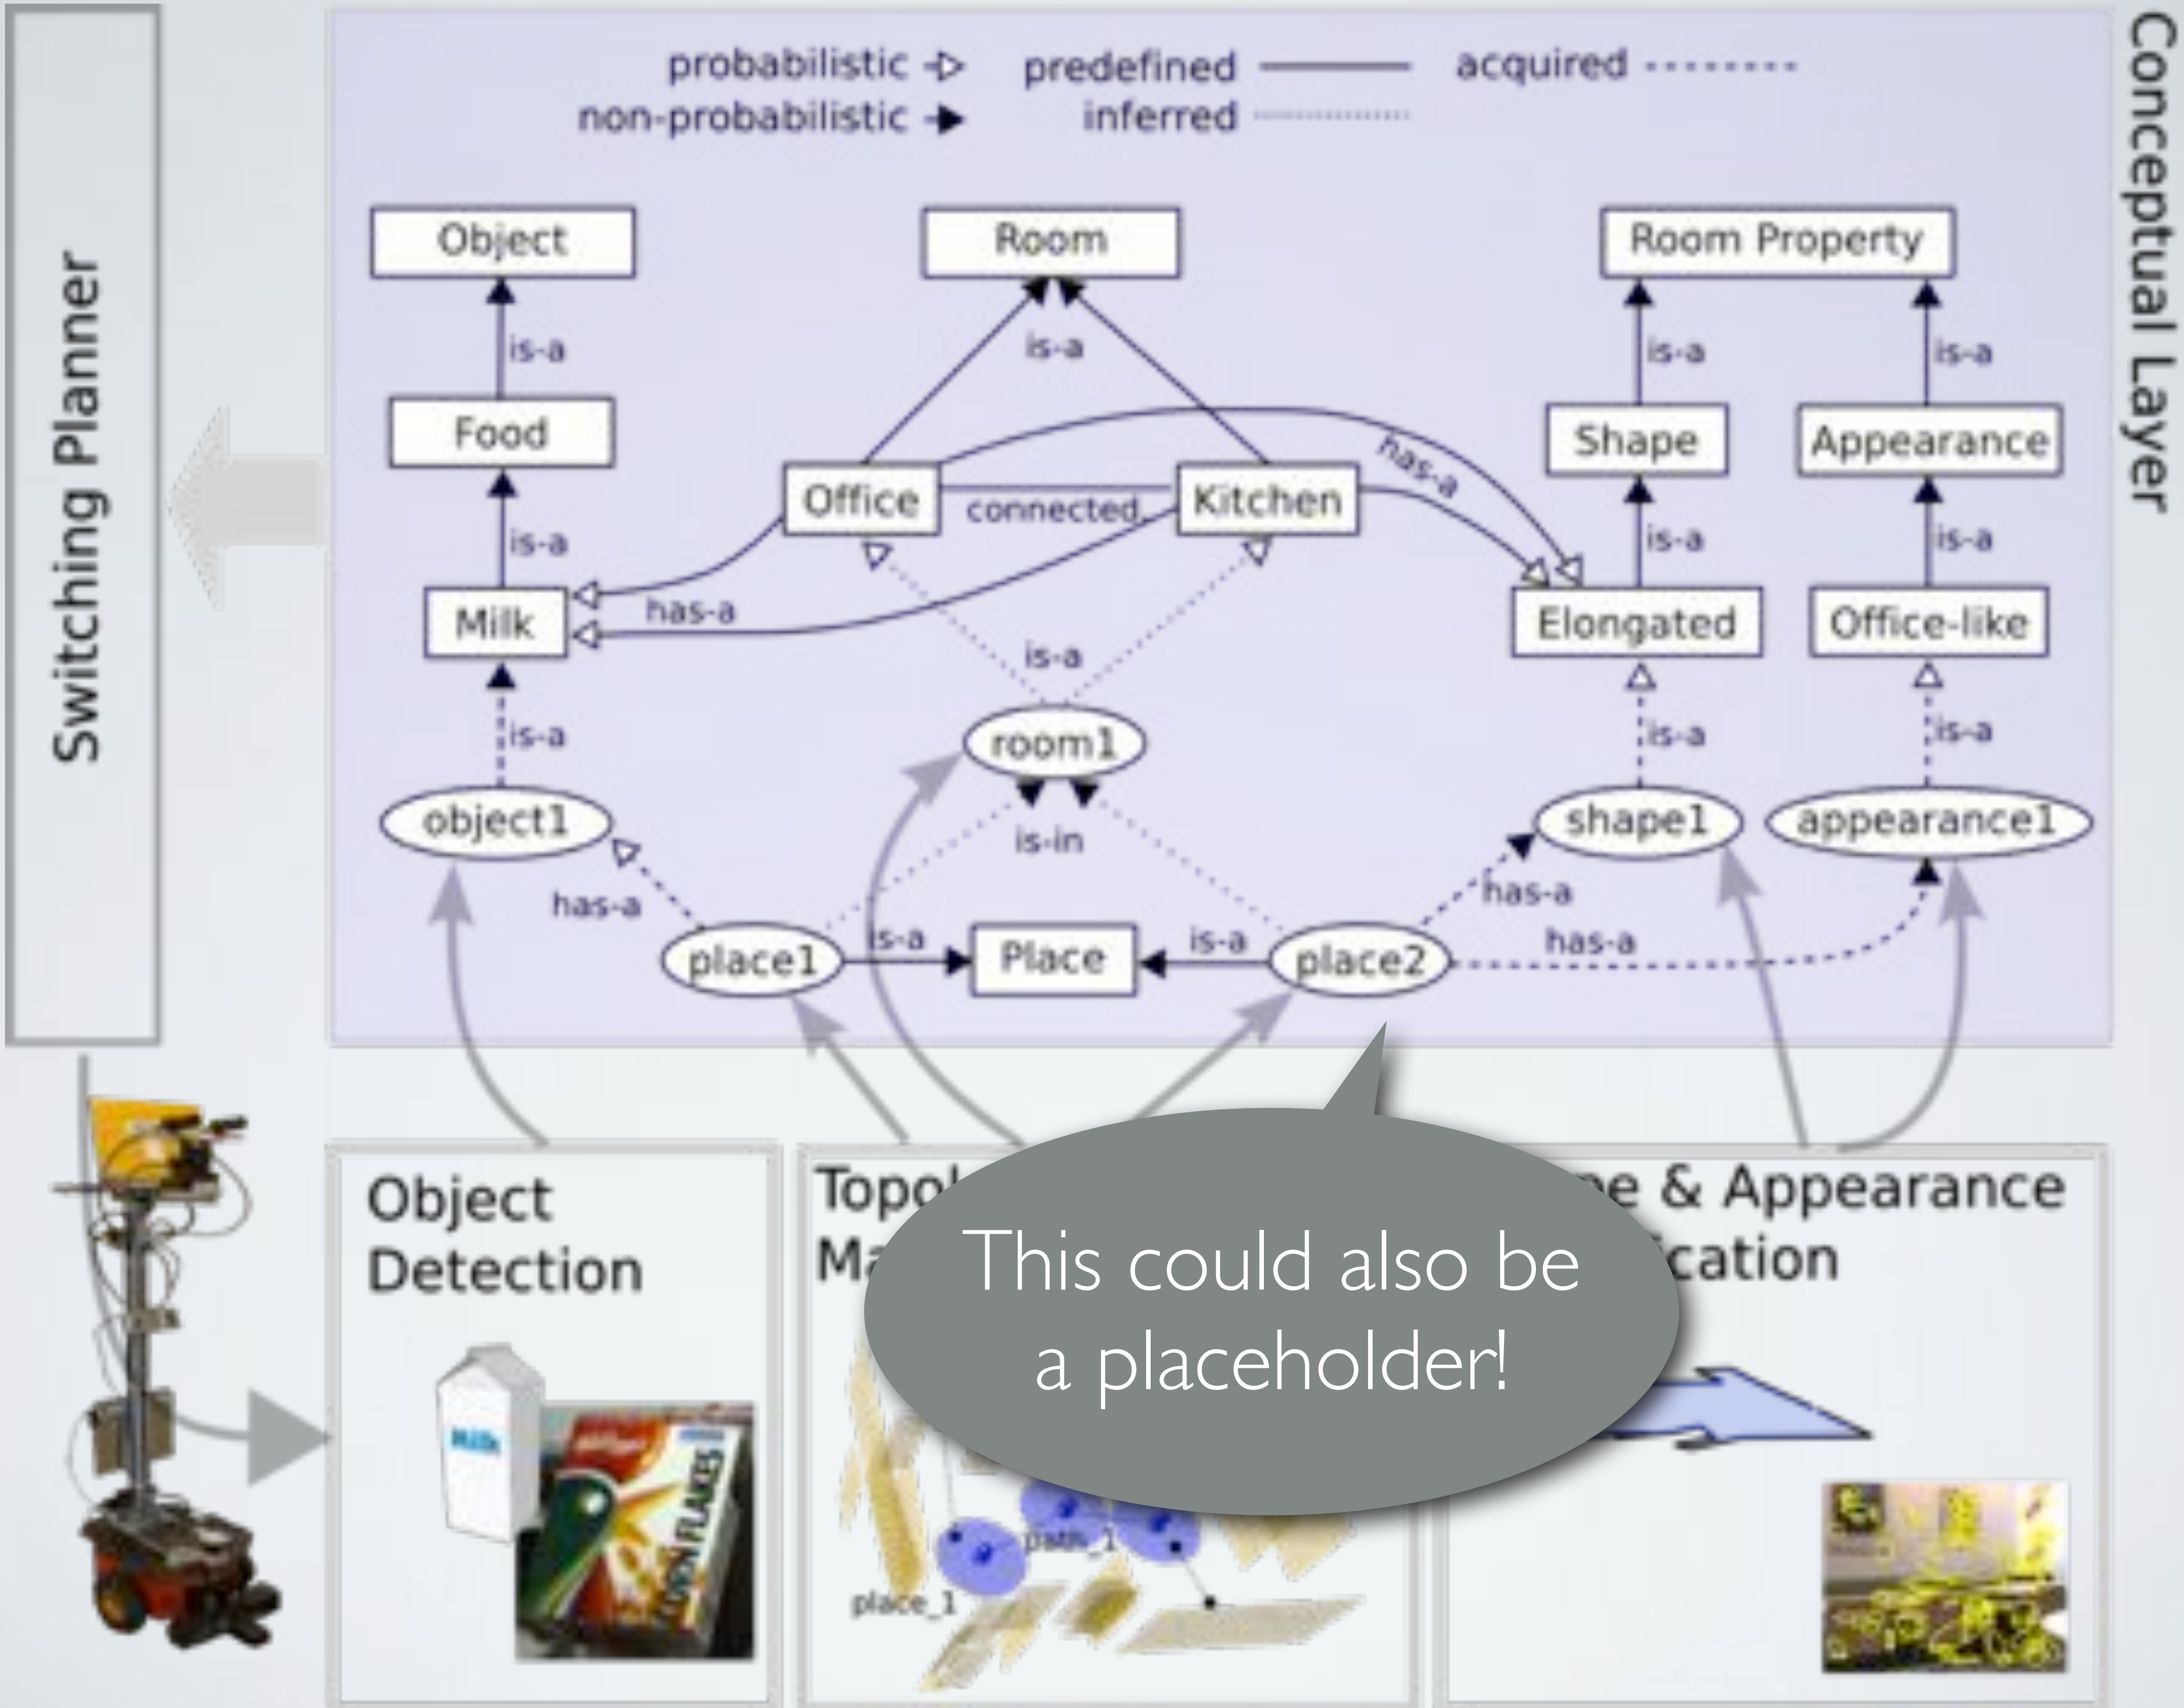

Switching Planner
Conceptual Layer
probabilistic
non-probabilistic
predefined
inferred
acquired
Object
Room
Room Property
Food
Office
connected
Kitchen
Shape
Appearance
Milk
Elongated
Office-like
room1
object1
shape1
appearance1
place1
Place
place2
is-a
has-a
is-in
Object Detection
path_1
place_1
This could also be a placeholder!

```
(exists
  (?o - visualobject)
  (and
    (= (label ?o) magazine)
    (kval [ROBOT] (related-to ?o))
  )
)
```

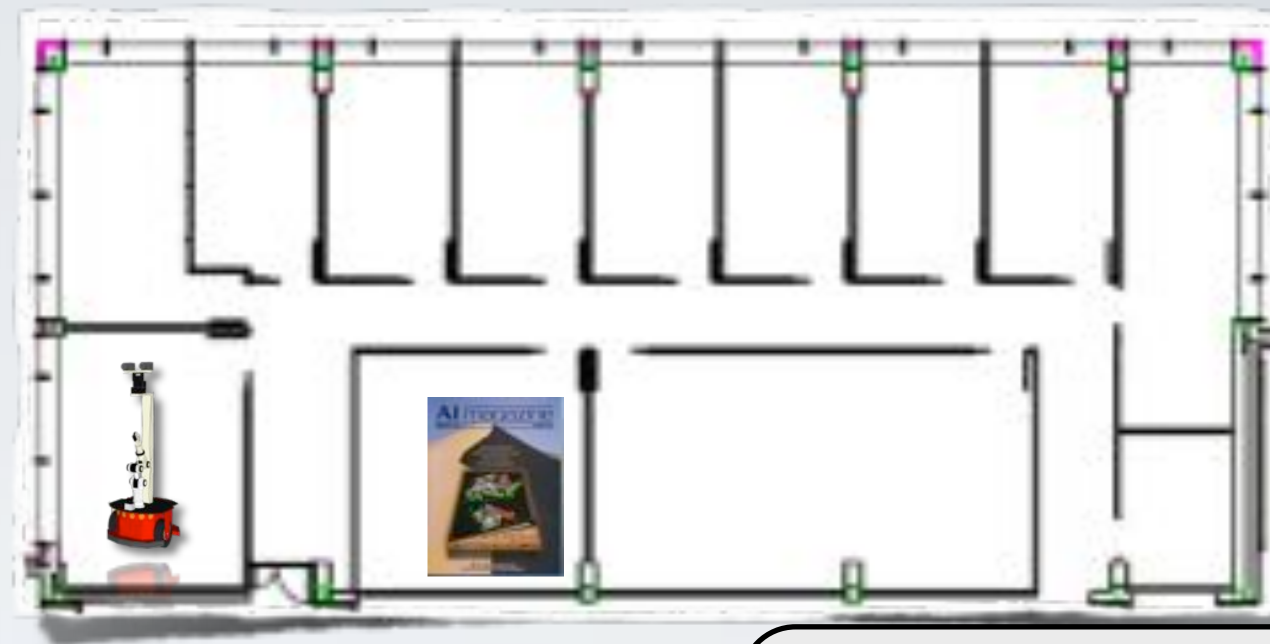

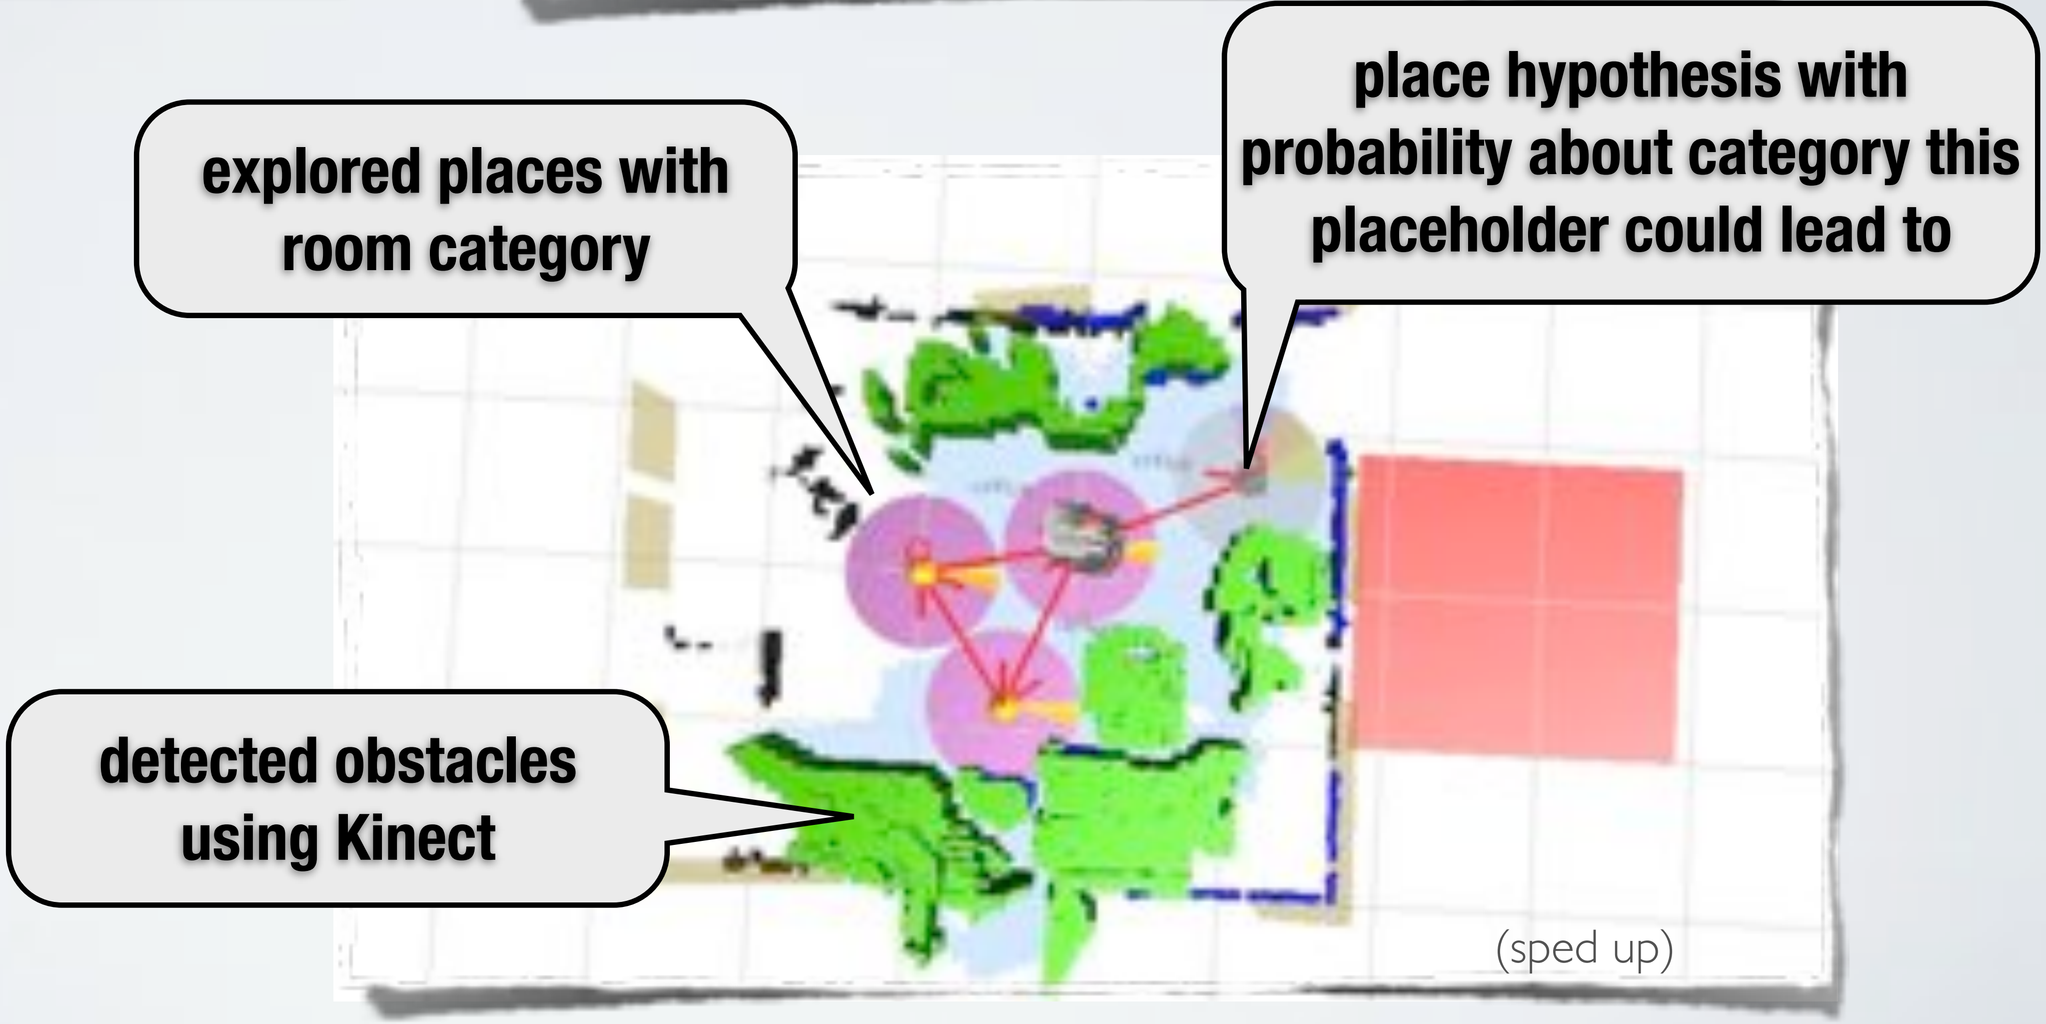

TASK-DRIVEN OPEN-WORLD

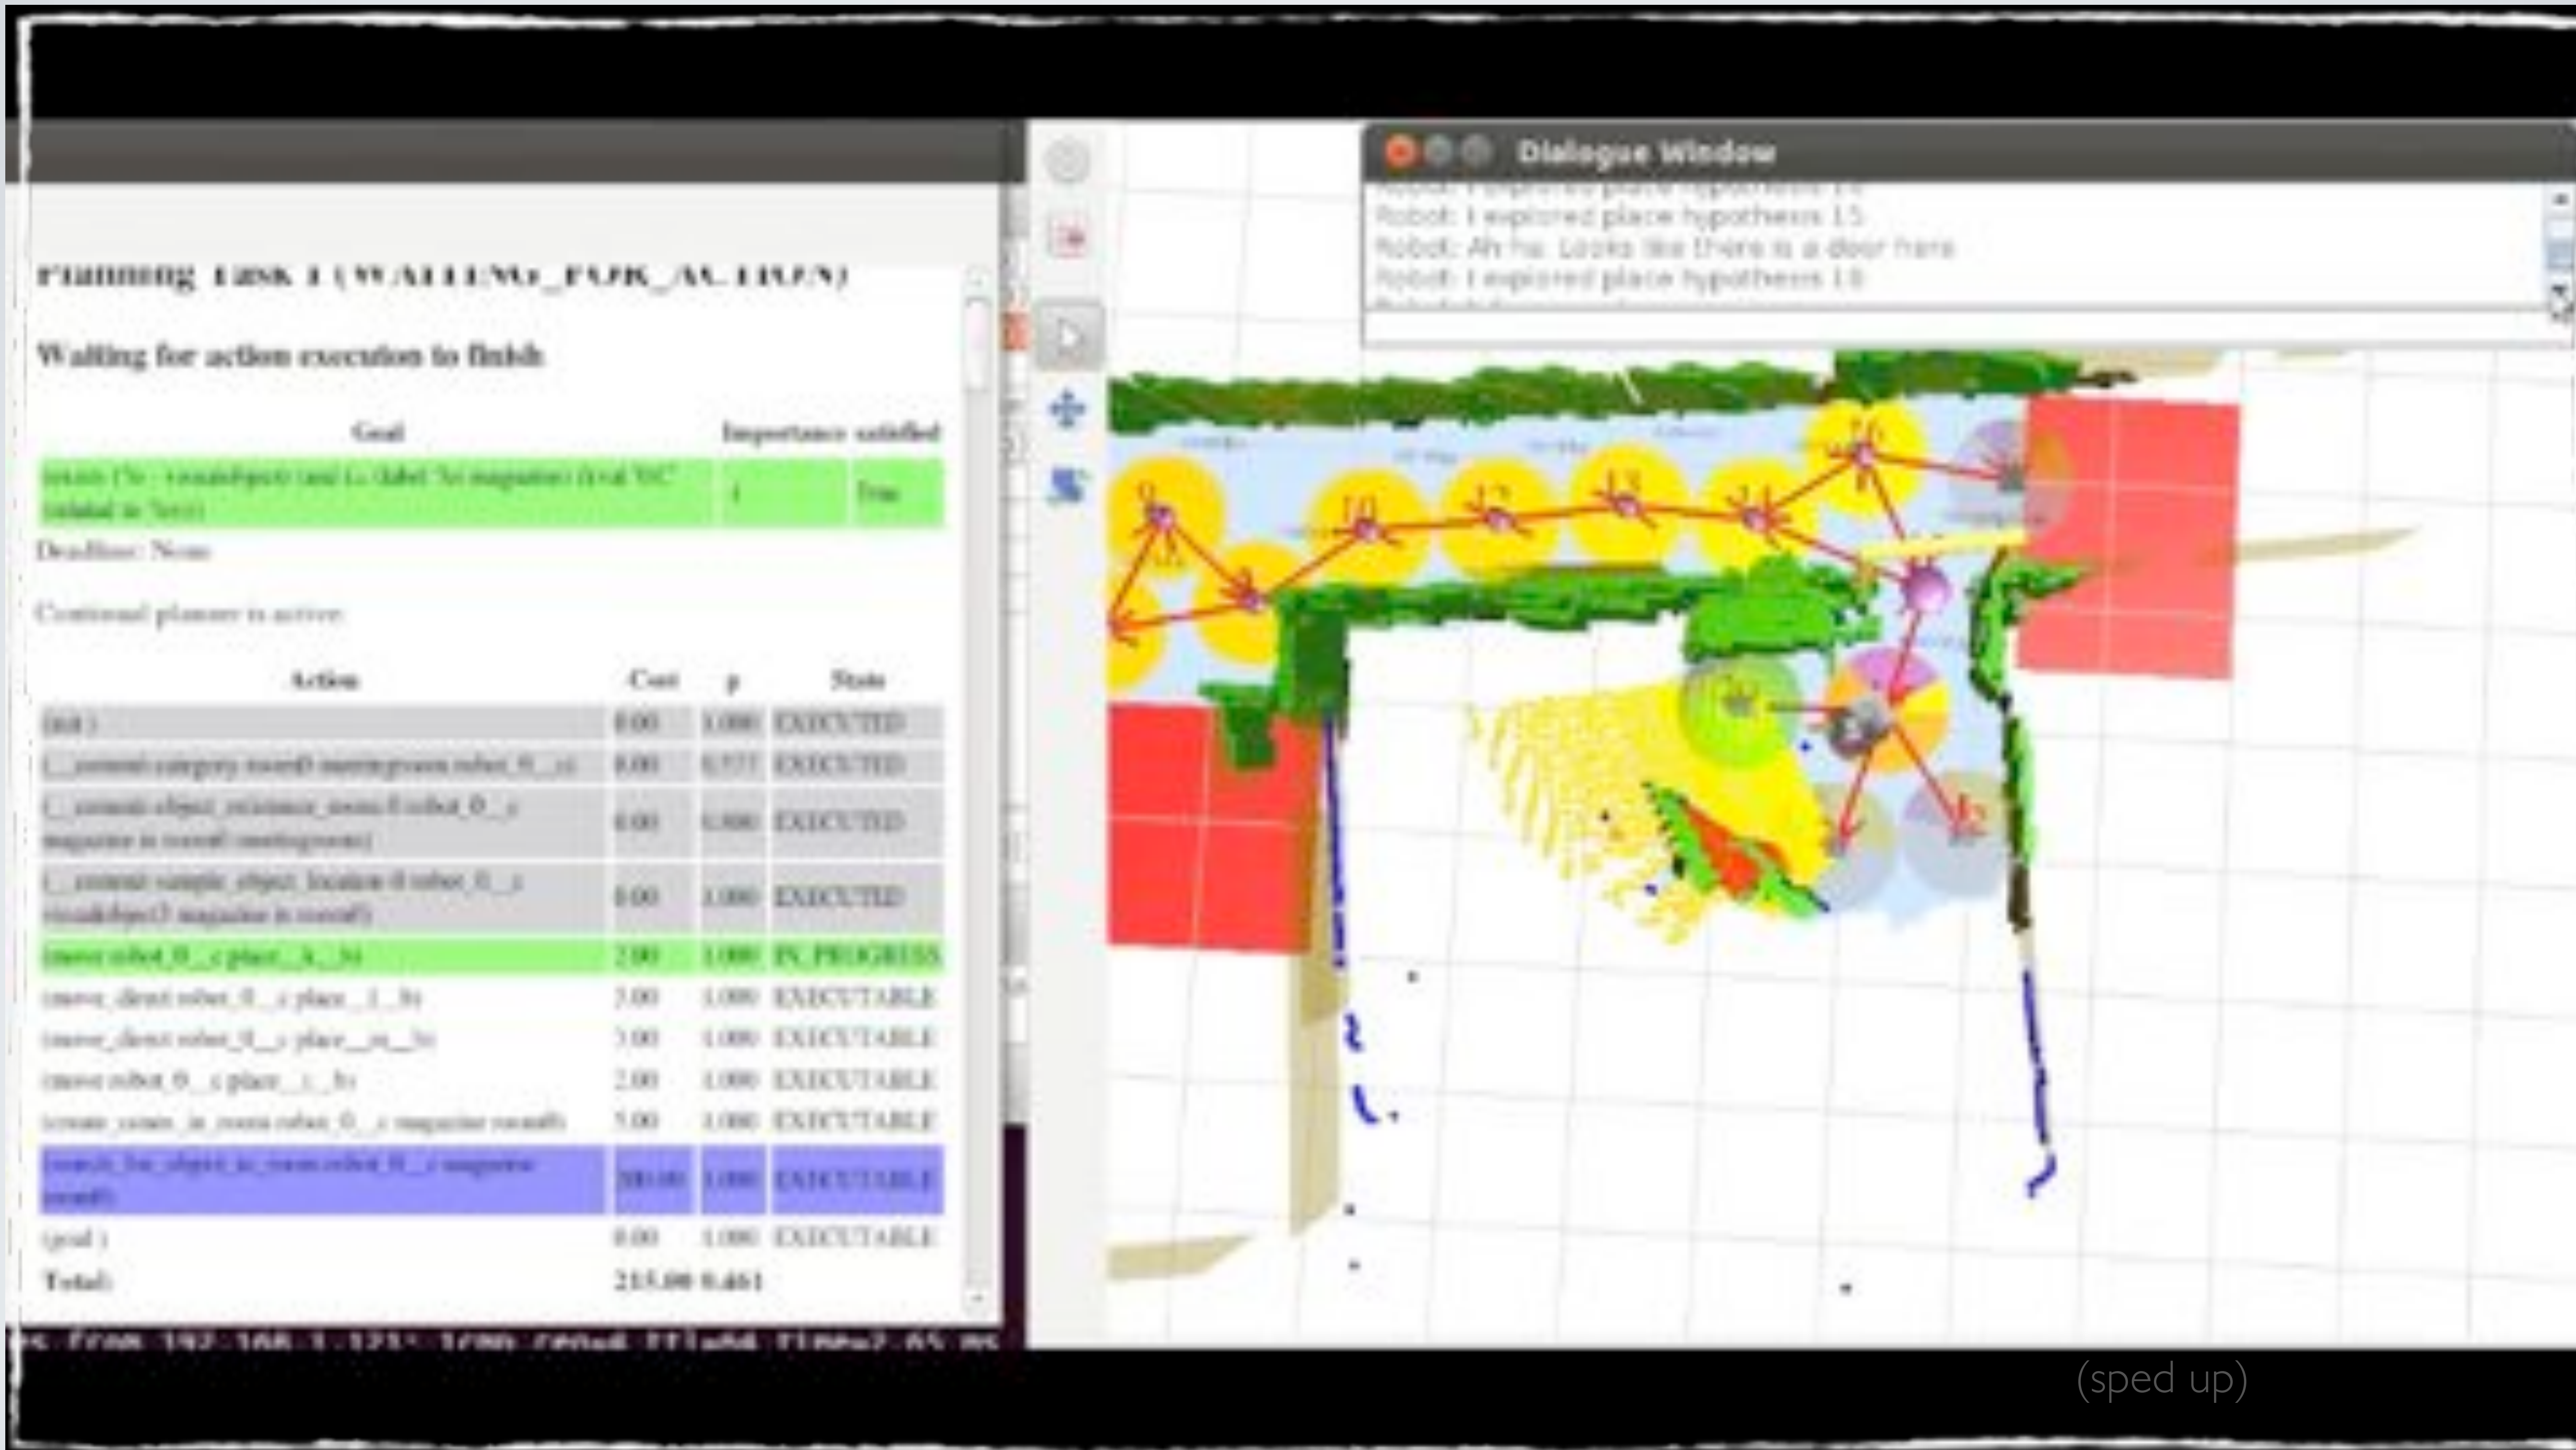

(sped up)

IF... THEN!

e object is not at (X,Y,Z)? ✓

Autonomous Visual Search

... the map would be unknown (no global coordinates)? ✓

Autonomous Exploration

... sensing would be uncertain? ✓

Probabilistic Planning

... the robot wouldn't know where the shop/ kitchen is? ✓

Room Categorisation

... the common-sense knowledge would be wrong? (✓)

Interactive Hypothesis Verification

Hanheide, M. et al., 2011. Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour. In Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI).

"Why did I not find the cornflakes?...

- Could that cornflakes are in a cupboard?
- Could it be that this room isn't a kitchen?
- Could it be that cornflakes are usually in offices?
- Could it be that the cornflakes are not in this room?"

EXPLAINING SURPRISES

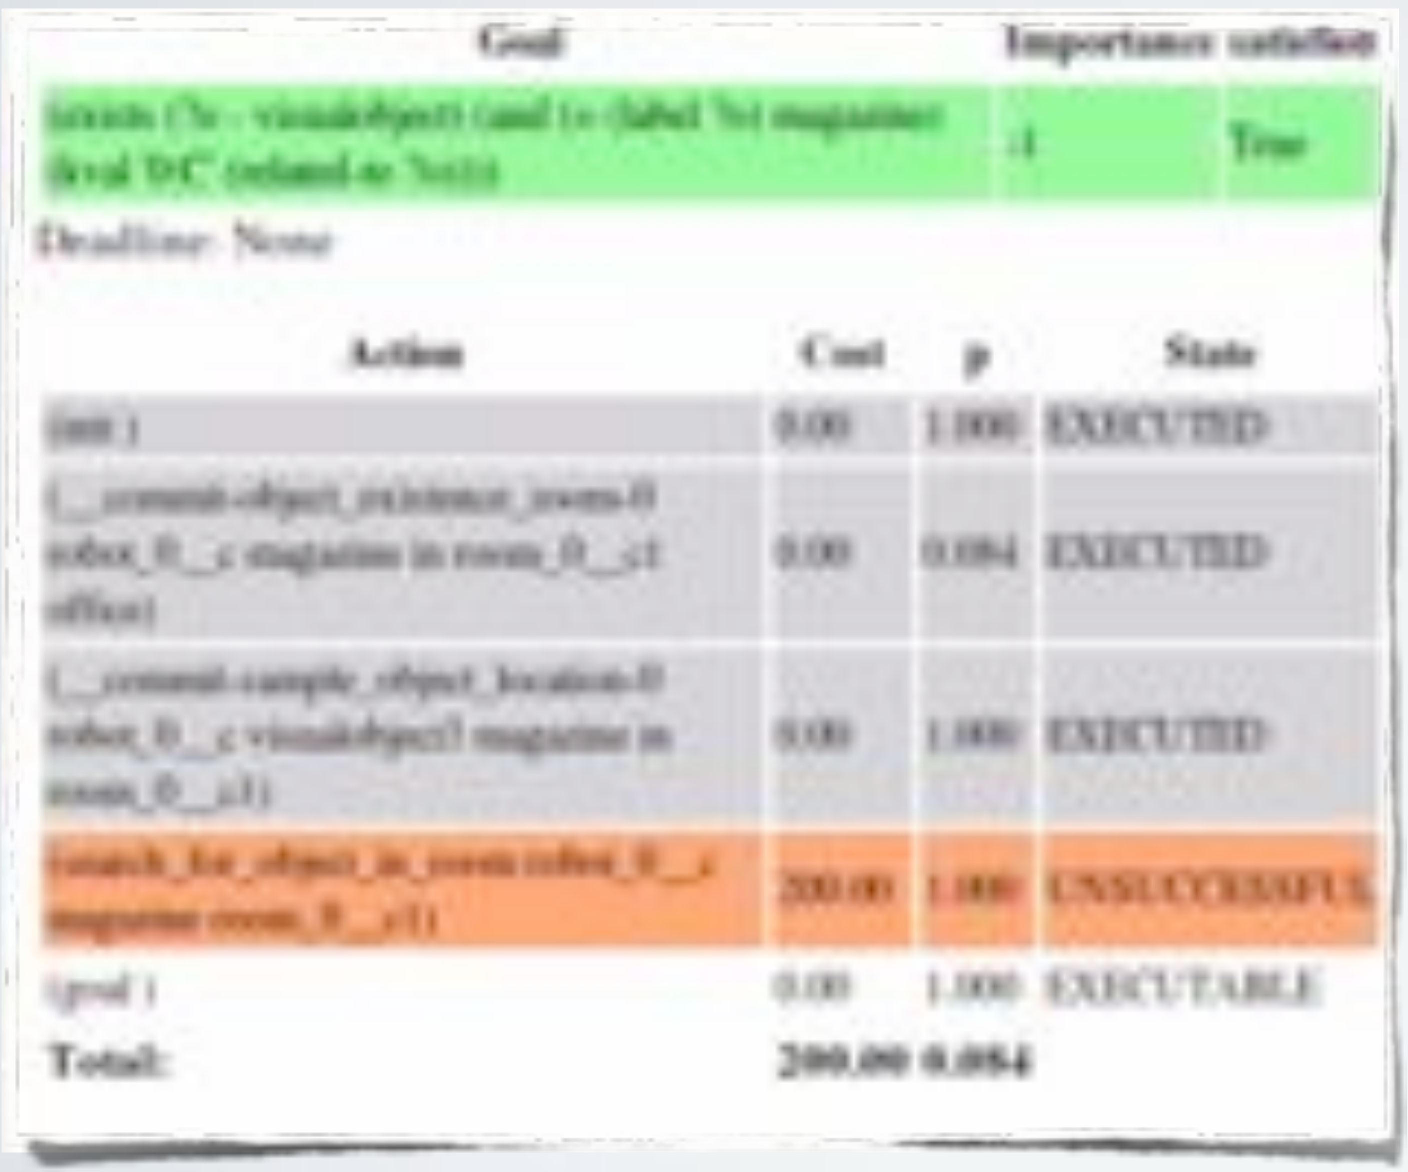

Goal	Importance	satisfied
[illegible]	-1	True

Deadline: None

Action	Cost	p	State
(init)	0.00	1.000	EXECUTED
[illegible]	0.00	0.084	EXECUTED
[illegible]	0.00	1.000	EXECUTED
[illegible]	200.00	1.000	UNSUCCESSFUL
(goal)	0.00	1.000	EXECUTABLE
Total:	200.00	0.084	

CREATING HYPOTHESES

Action	Cost	p	State
(init)	0.00	1.000	EXECUTED
(__commit-object_existence_room-0 robot_0__c magazine in room_0__c1 office)	0.00	0.084	EXECUTED
(__commit-sample_object_location-0 robot_0__c visualobject1 magazine in room_0__c1)	0.00	1.000	EXECUTED
(search_for_object_in_room robot_0__c magazine room_0__c1)	200.00	1.000	UNSUCCESSFUL
(goal)	0.00	1.000	EXECUTABLE
Total:	200.00	0.084	

new assumptions

Goal
(and (kval '0:C' (dora__inobject magazine container meetingroom)))
(kval '0:C' (entity-exists '6:y'))
(and (kval '0:C' (dora__inroom container meetingroom)))
(kval '0:C' (entity-exists '7:y'))

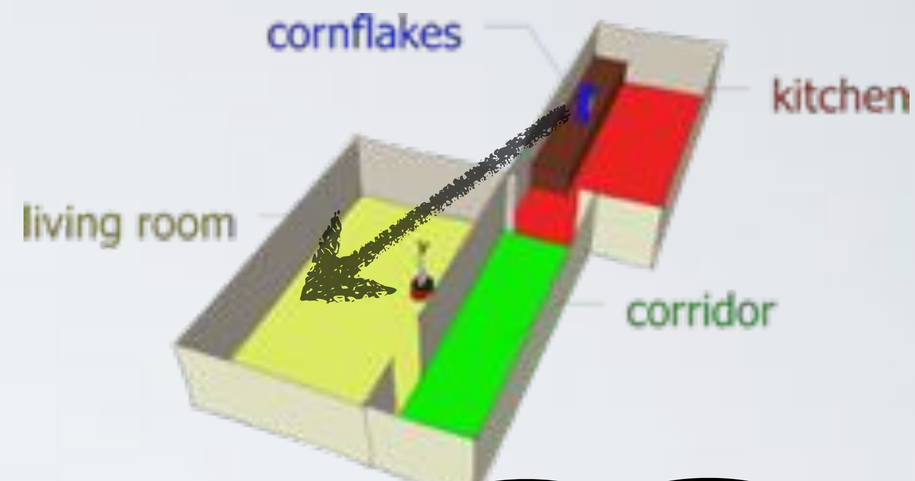

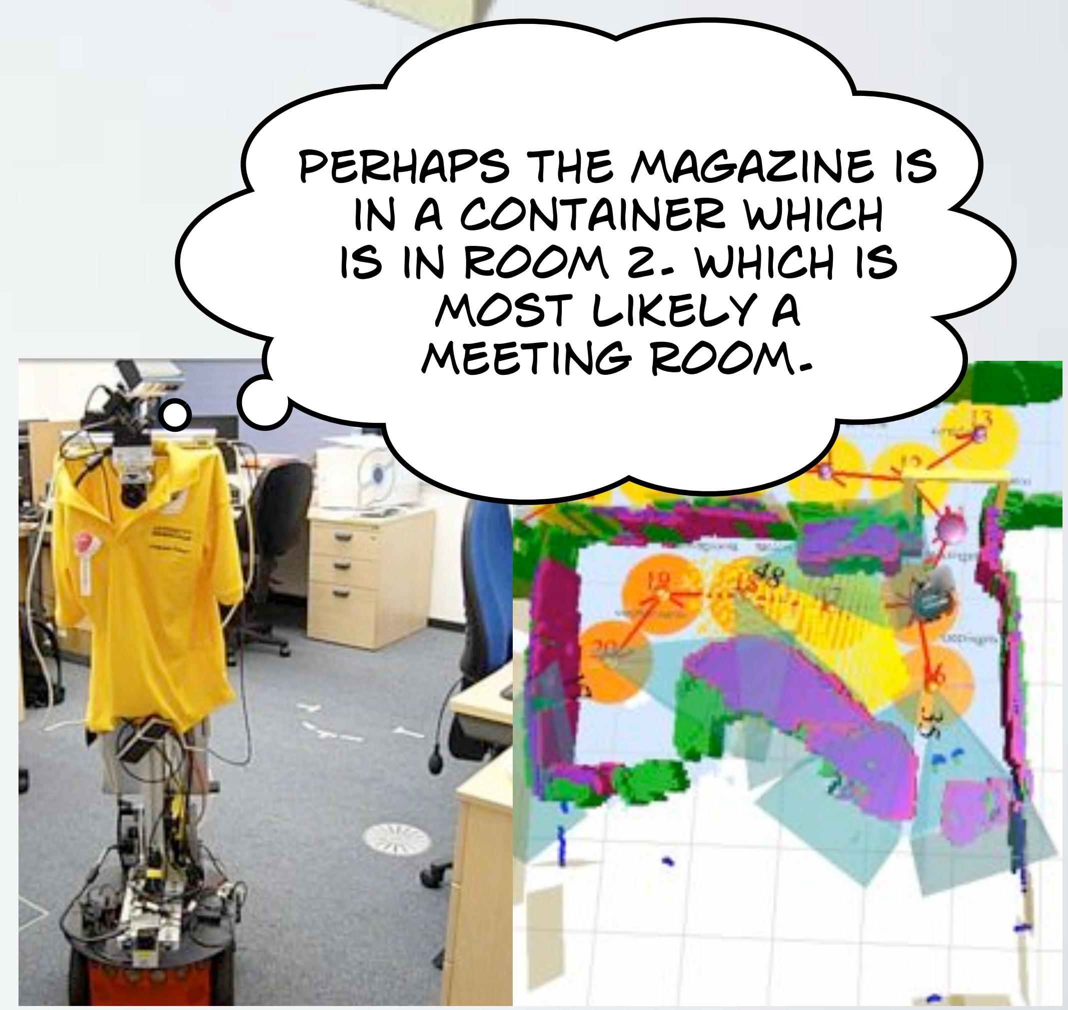

VERIFYING HYPOTHESES

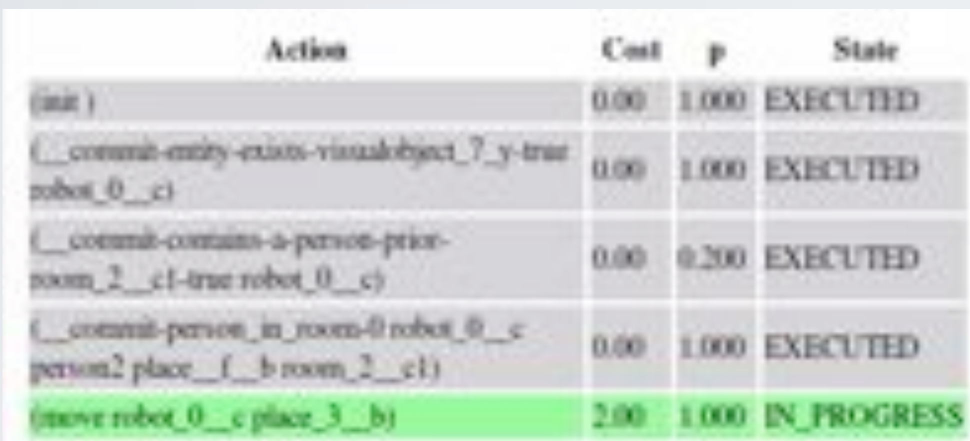

Action	Cost	p	State
(init)	0.00	1.000	EXECUTED
(__commit-entity-exists-visualobject_7_y-true robot_0__c)	0.00	1.000	EXECUTED
(__commit-contains-a-person-prior-room_2__c1-true robot_0__c)	0.00	0.200	EXECUTED
(__commit-person_in_room-0 robot_0__c person2 place__f__b room_2__c1)	0.00	1.000	EXECUTED
(move robot_0__c place_3__b)	2.00	1.000	IN_PROGRESS

verification plan

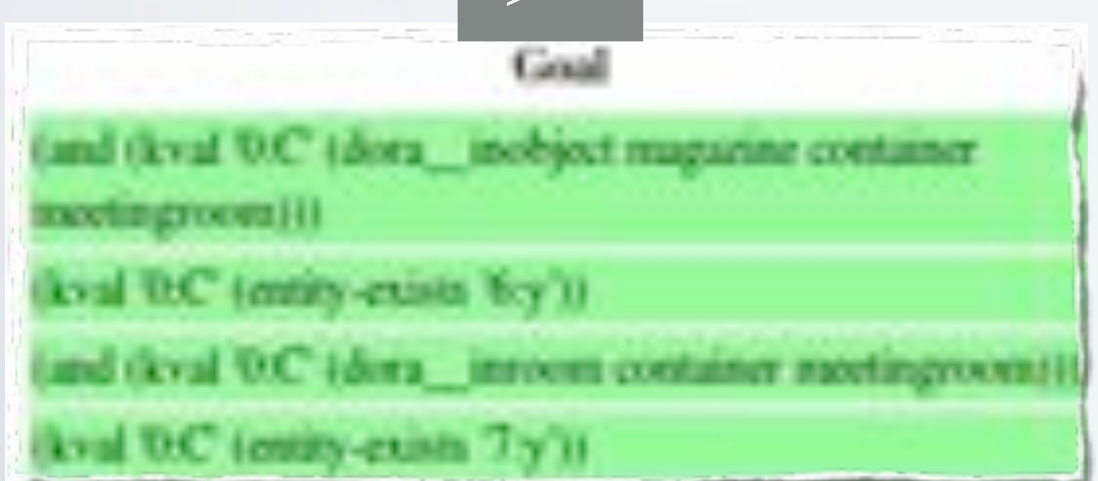

Goal
(and (kval '0:C' (dora__inobject magazine container meetingroom)))
(kval '0:C' (entity-exists '6:y'))
(and (kval '0:C' (dora__inroom container meetingroom)))
(kval '0:C' (entity-exists '7:y'))

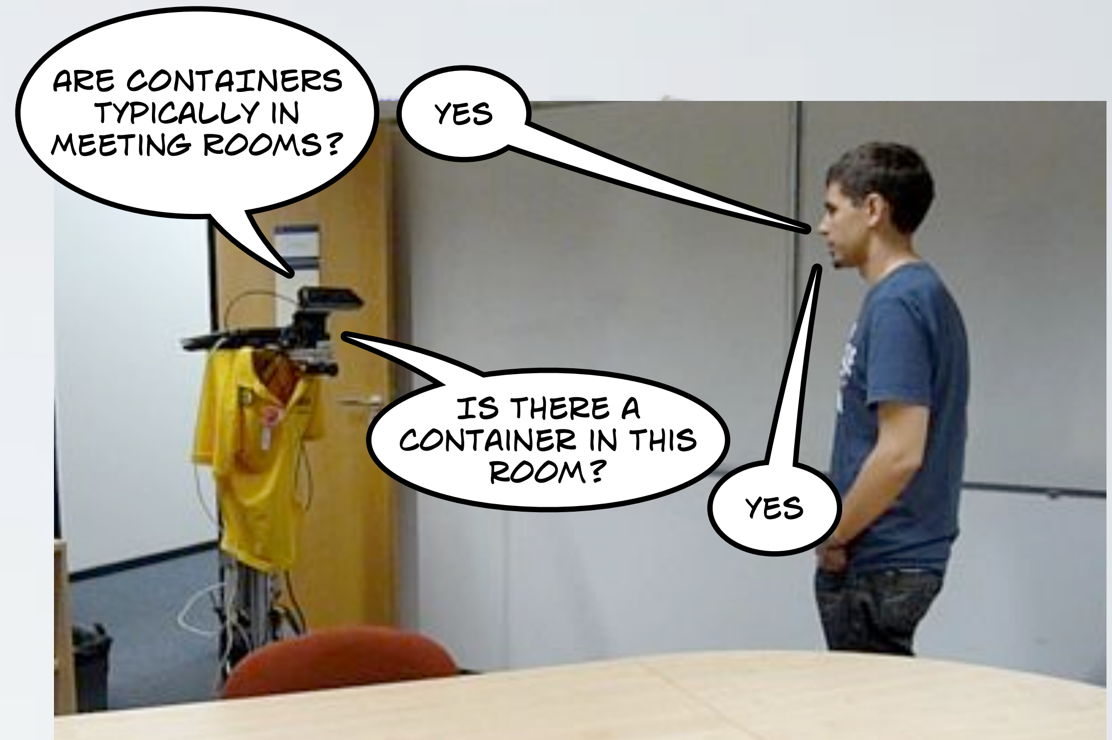

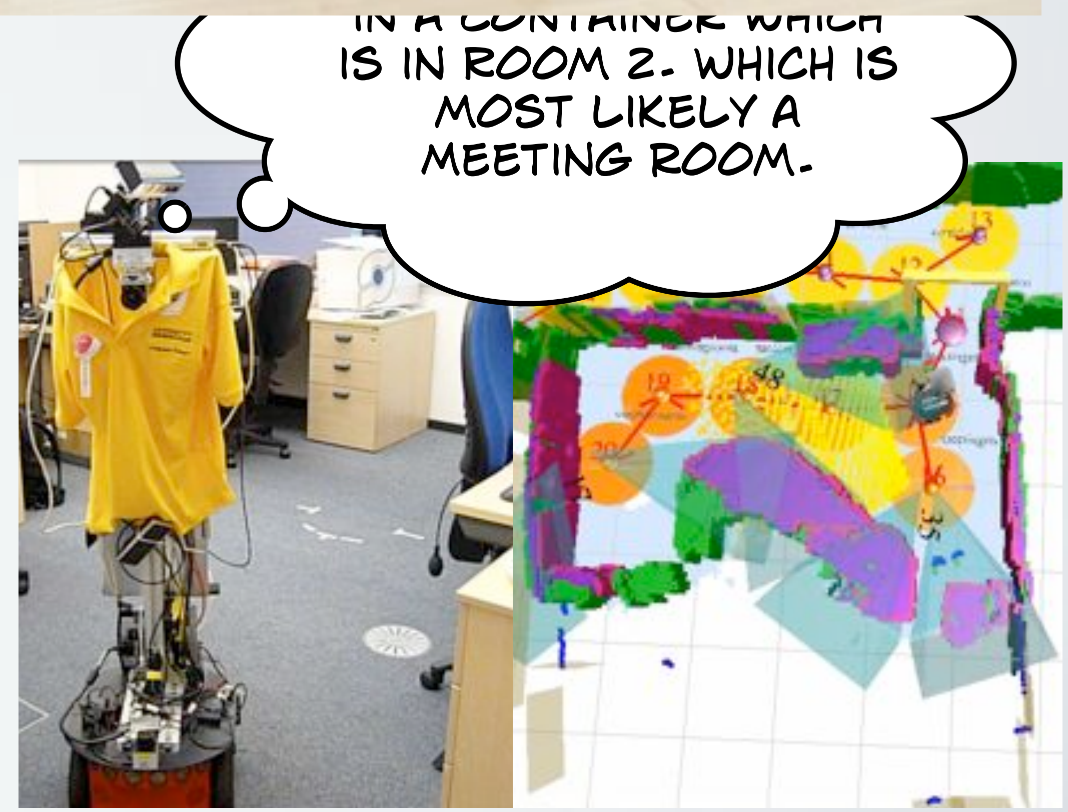

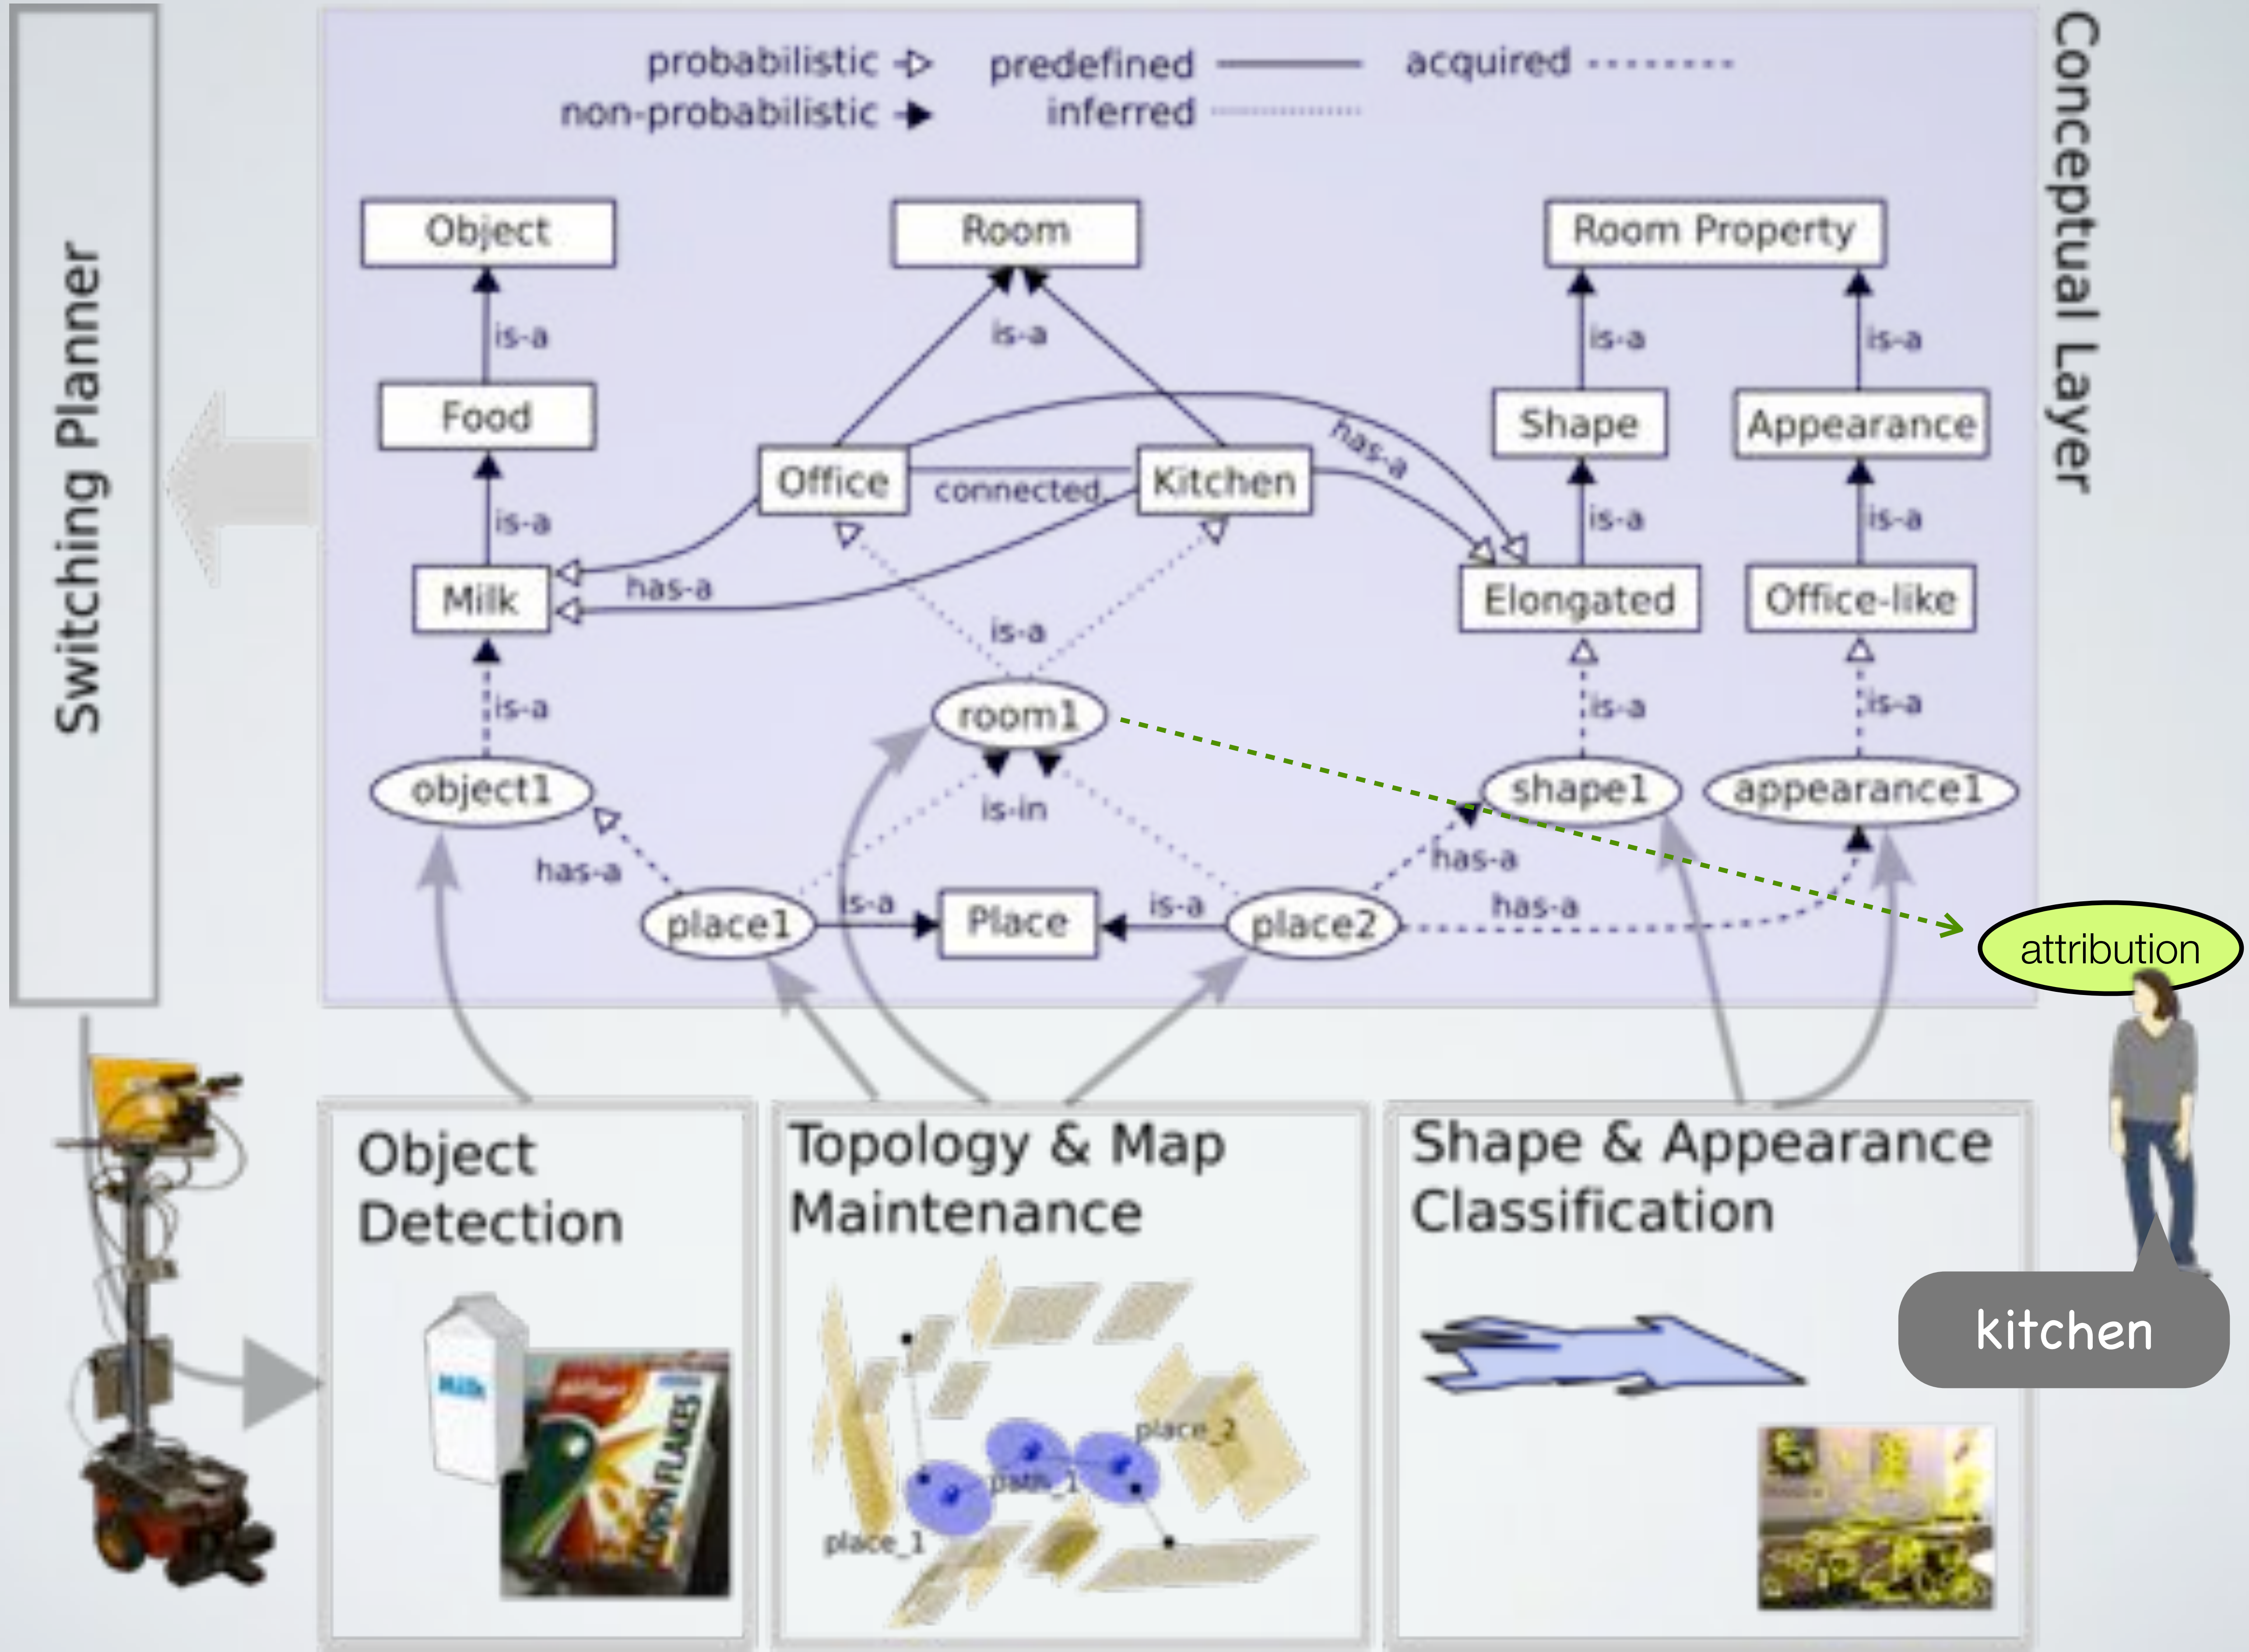

Switching Planner
Conceptual Layer
probabilistic
non-probabilistic
predefined
inferred
acquired
Object
Room
Room Property
is-a
Food
Office
connected
Kitchen
has-a
Shape
Appearance
Milk
Elongated
Office-like
room1
object1
is-in
shape1
appearance1
place1
Place
place2
attribution
Object Detection
Topology & Map Maintenance
Shape & Appearance Classification
place_1
path_1
place_2
kitchen

"IS IT AN OFFICE?"

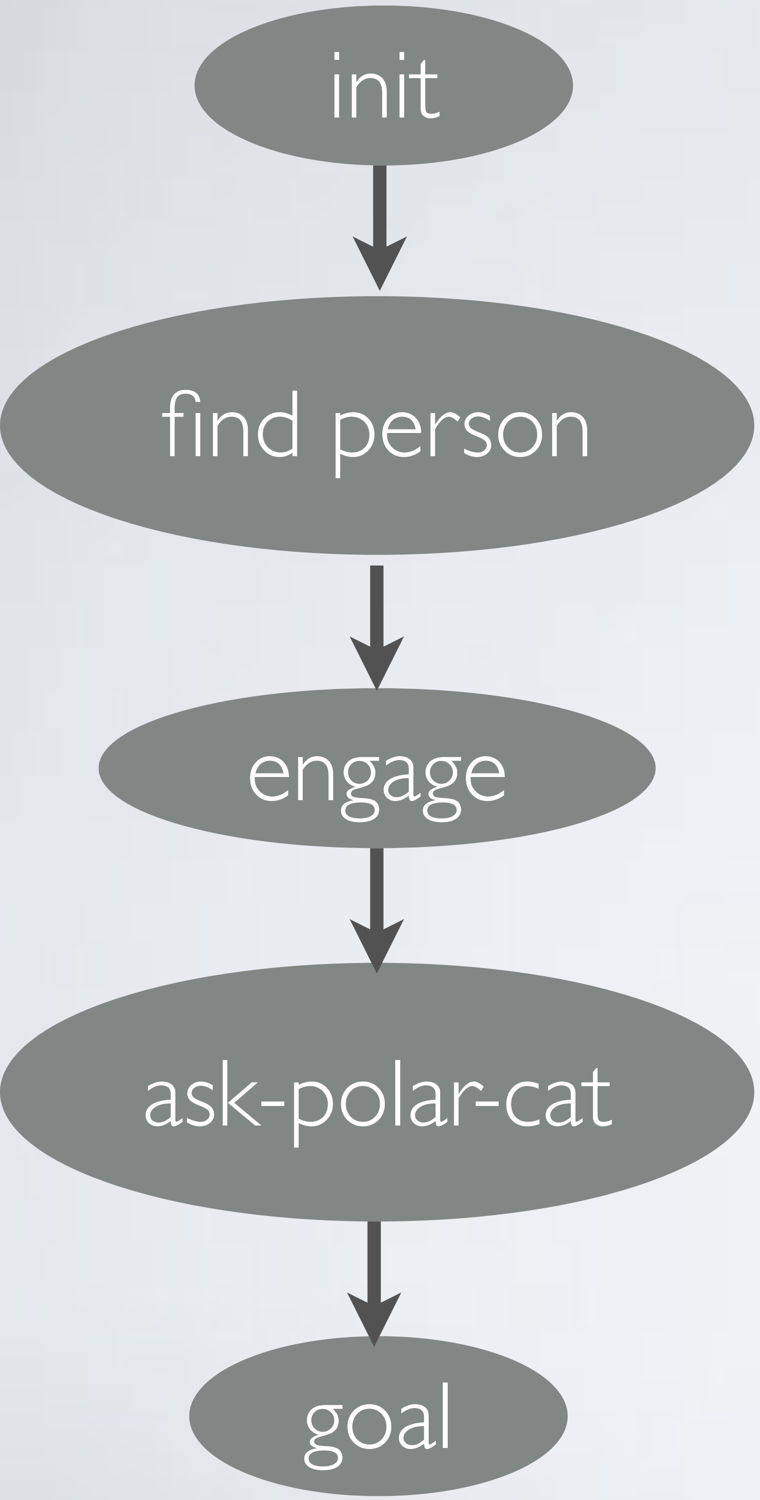

IF... THEN!

e object is not at (X,Y,Z)? ✓

Autonomous Visual Search

... the map would be unknown (no global coordinates)?

Autonomous Exploration

... sensing would be uncertain? ✓

Probabilistic Planning

... the robot wouldn't know where the shop/kitchen is? ✓

Room Categorisation

... the common-sense knowledge would be wrong? ✓

Interactive Hypothesis Verification

BUT...

I can't
get you a
sandwich...

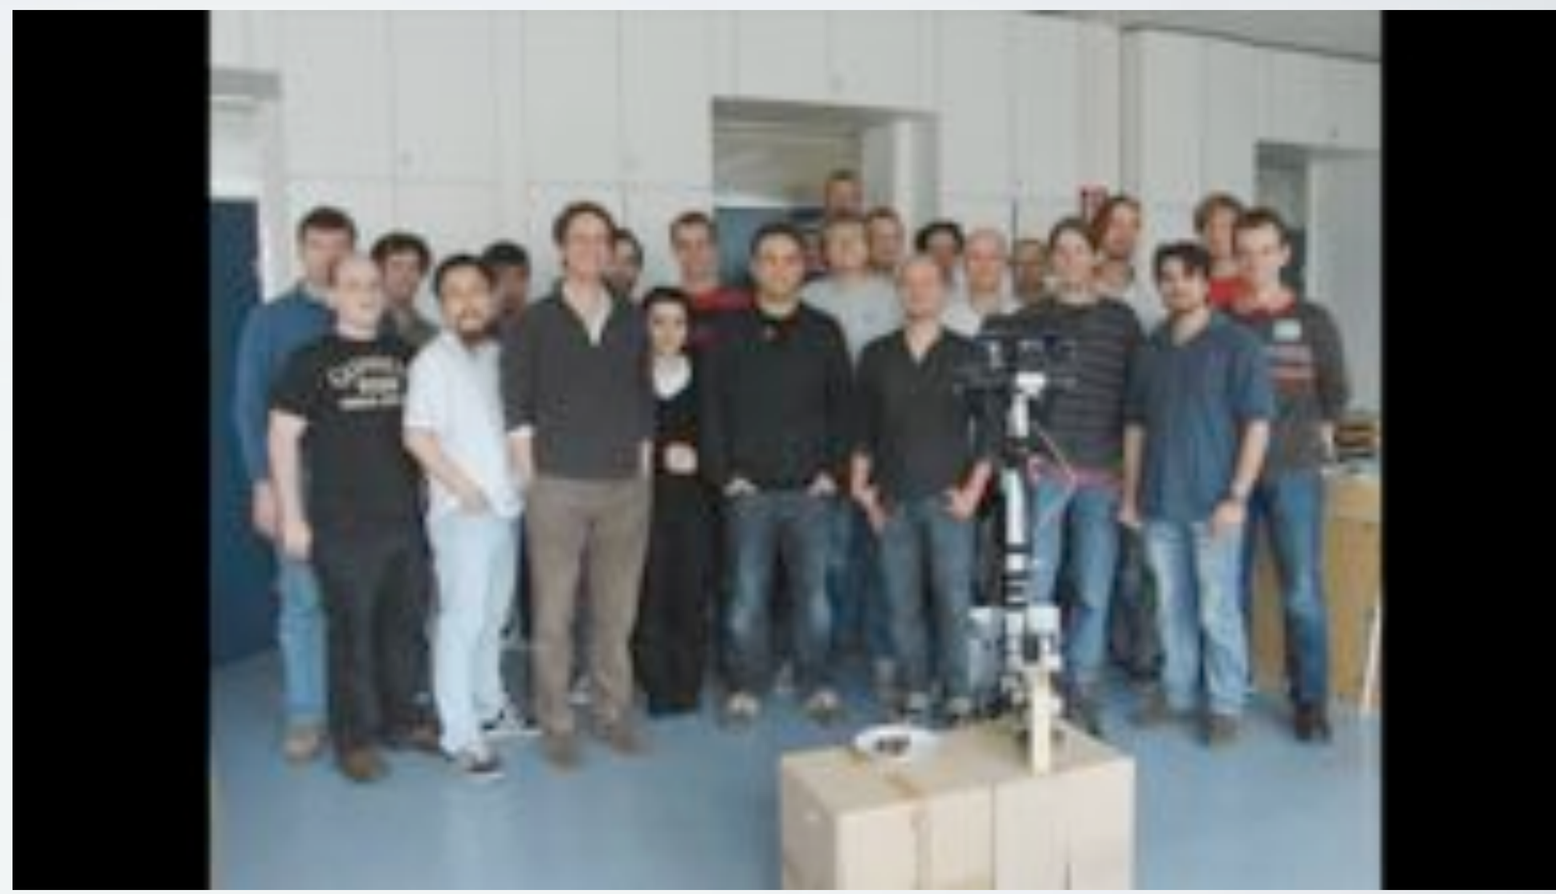

THE END

- Dora: A system (not only) for object search
  - exploiting common-sense knowledge for *efficiency*
  - deals with uncertainty for *robustness*
  - explores its environment and learns through interaction to *extend* its knowledge

Grazie!

provide *services*...

... in *human-shared* environments

operate in the *real* world...

... even though it's full of *uncertainties*

be *adaptive*...

... but don't learn from *scratch*

exploit common-sense!

deal with uncertainties!

http://cogx.eu